

Building Query Compilers
(Under Construction)
[expected time to completion: 5 years]

Guido Moerkotte

September 3, 2009

Contents

I	Basics	3
1	Introduction	5
1.1	General Remarks	5
1.2	DBMS Architecture	5
1.3	Interpretation versus Compilation	6
1.4	Requirements for a Query Compiler	9
1.5	Search Space	11
1.6	Generation versus Transformation	12
1.7	Focus	12
1.8	Organization of the Book	13
2	Textbook Query Optimization	15
2.1	Example Query and Outline	15
2.2	Algebra	16
2.3	Canonical Translation	17
2.4	Logical Query Optimization	20
2.5	Physical Query Optimization	24
2.6	Discussion	25
3	Join Ordering	31
3.1	Queries Considered	31
3.1.1	Query Graph	32
3.1.2	Join Tree	33
3.1.3	Simple Cost Functions	34
3.1.4	Classification of Join Ordering Problems	40
3.1.5	Search Space Sizes	41
3.1.6	Problem Complexity	45
3.2	Deterministic Algorithms	47
3.2.1	Heuristics	47
3.2.2	Determining the Optimal Join Order in Polynomial Time	49
3.2.3	The Maximum-Value-Precedence Algorithm	56
3.2.4	Dynamic Programming	61
3.2.5	Memoization	78
3.2.6	Join Ordering by Generating Permutations	79
3.2.7	A Dynamic Programming based Heuristics for Chain Queries	81
3.2.8	Transformation-Based Approaches	94

3.3	Probabilistic Algorithms	101
3.3.1	Generating Random Left-Deep Join Trees with Cross Products	101
3.3.2	Generating Random Join Trees with Cross Products	103
3.3.3	Generating Random Join Trees without Cross Products	107
3.3.4	Quick Pick	116
3.3.5	Iterative Improvement	116
3.3.6	Simulated Annealing	117
3.3.7	Tabu Search	118
3.3.8	Genetic Algorithms	119
3.4	Hybrid Algorithms	122
3.4.1	Two Phase Optimization	122
3.4.2	AB-Algorithm	122
3.4.3	Toured Simulated Annealing	122
3.4.4	GOO-II	123
3.4.5	Iterative Dynamic Programming	123
3.5	Ordering Order-Preserving Joins	123
3.6	Characterizing Search Spaces	131
3.6.1	Complexity Thresholds	131
3.7	Discussion	135
3.8	Bibliography	136
4	Database Items, Building Blocks, and Access Paths	137
4.1	Disk Drive	137
4.2	Database Buffer	146
4.3	Physical Database Organization	147
4.4	Slotted Page and Tuple Identifier (TID)	150
4.5	Physical Record Layouts	151
4.6	Physical Algebra (Iterator Concept)	152
4.7	Simple Scan	152
4.8	Scan and Attribute Access	153
4.9	Temporal Relations	155
4.10	Table Functions	155
4.11	Indexes	156
4.12	Single Index Access Path	158
4.12.1	Simple Key, No Data Attributes	158
4.12.2	Complex Keys and Data Attributes	163
4.13	Multi Index Access Path	165
4.14	Indexes and Joins	167
4.15	Remarks on Access Path Generation	172
4.16	Counting the Number of Accesses	173
4.16.1	Counting the Number of Direct Accesses	173
4.16.2	Counting the Number of Sequential Accesses	182
4.16.3	Pointers into the Literature	187
4.17	Disk Drive Costs for N Uniform Accesses	188
4.17.1	Number of Qualifying Cylinders, Tracks, and Sectors	188
4.17.2	Command Costs	189

4.17.3	Seek Costs	189
4.17.4	Settle Costs	191
4.17.5	Rotational Delay Costs	191
4.17.6	Head Switch Costs	193
4.17.7	Discussion	193
4.18	Concluding Remarks	194
4.19	Bibliography	194
 II Foundations		197
5	Logic, Null, and Boolean Expressions	199
5.1	Two-valued logic	199
5.2	NULL values and three valued logic	199
5.3	Simplifying Boolean Expressions	200
5.4	Optimizing Boolean Expressions	200
5.5	Bibliography	200
6	Functional Dependencies	203
6.1	Functional Dependencies	203
6.2	Functional Dependencies in the presence of NULL values	204
6.3	Deriving Functional Dependencies over algebraic operators	204
6.4	Bibliography	204
7	An Algebra for Sets, Bags, and Sequences	205
7.1	Sets, Bags, and Sequences	205
7.1.1	Sets	205
7.1.2	Duplicate Data: Bags	206
7.1.3	Ordered Data: Sequences	208
7.2	Algebra	209
7.2.1	The Operators	210
7.2.2	Preliminaries	211
7.2.3	Operator Signatures	213
7.2.4	Selection	215
7.2.5	Projection	215
7.2.6	Map	215
7.2.7	Join Operators	215
7.2.8	Linearity and Reorderability	217
7.2.9	Reordering of joins and outer-joins	220
7.2.10	Basic Equivalences for d-Join and Grouping	223
7.2.11	Simplifying Expressions Containing Joins	224
7.2.12	Reordering Joins and Grouping	226
7.2.13	ToDo	233
7.3	Logical Algebra for Sequences	233
7.3.1	Introduction	233
7.3.2	Algebraic Operators	235
7.3.3	Equivalences	238

7.3.4	Bibliography	238
7.4	Literature	238
8	Declarative Query Representation	241
8.1	Calculus Representations	241
8.2	Datalog	241
8.3	Tableaux Representation	241
8.4	Monoid Comprehension	241
8.5	Expressiveness	241
8.6	Bibliography	241
9	Translation and Lifting	243
9.1	Query Language to Calculus	243
9.2	Query Language to Algebra	243
9.3	Calculus to Algebra	243
9.4	Algebra to Calculus	243
9.5	Bibliography	243
10	Query Equivalence, Containment, Minimization, and Factorization	245
10.1	Set Semantics	246
10.1.1	Conjunctive Queries	246
10.1.2	... with Inequalities	248
10.1.3	... with Negation	249
10.1.4	... under Constraints	249
10.1.5	... with Aggregation	249
10.2	Bag Semantics	249
10.2.1	Conjunctive Queries	249
10.3	Sequences	250
10.3.1	Path Expressions	250
10.4	Minimization	251
10.5	Detecting common subexpressions	251
10.5.1	Simple Expressions	251
10.5.2	Algebraic Expressions	251
10.6	Bibliography	251
III	Rewrite Techniques	253
11	Simple Rewrites	255
11.1	Simple Adjustments	255
11.1.1	Rewriting Simple Expressions	255
11.1.2	Normal forms for queries with disjunction	257
11.2	Deriving new predicates	257
11.2.1	Collecting conjunctive predicates	257
11.2.2	Equality	257
11.2.3	Inequality	258

11.2.4	Aggregation	259
11.2.5	ToDo	261
11.3	Predicate Push-Down and Pull-Up	261
11.4	Eliminating Redundant Joins	261
11.5	Distinct Pull-Up and Push-Down	261
11.6	Set-Valued Attributes	261
11.6.1	Introduction	261
11.6.2	Preliminaries	262
11.6.3	Query Rewrite	263
11.7	Bibliography	264
12	View Merging	267
12.1	View Resolution	267
12.2	Simple View Merging	267
12.3	Predicate Move Around (Predicate pull-up and push-down) . . .	268
12.4	Complex View Merging	269
12.4.1	Views with Distinct	269
12.4.2	Views with Group-By and Aggregation	270
12.4.3	Views in IN predicates	271
12.4.4	Final Remarks	271
12.5	Bibliography	272
13	Quantifier treatment	273
13.1	Pseudo-Quantifiers	273
13.2	Existential quantifier	274
13.3	Universal quantifier	274
13.4	Bibliography	278
14	Unnesting Nested Queries	279
14.1	Classification of nested queries	279
14.2	Queries of Type A	280
14.3	Queries of Type N	281
14.4	Queries of Type J	284
14.5	Queries of Type JA	286
14.6	Alternative locations	287
14.7	Different Kinds of Dependency	289
14.8	Unnesting IN	291
14.9	Further reading	291
14.10	History	291
14.11	Bibliography	292
14.12	ToDo	292
15	Optimizing Queries with Materialized Views	295
15.1	Conjunctive Views	295
15.2	Views with Grouping and Aggregation	295
15.3	Views with Disjunction	295
15.4	Bibliography	295

16 Semantic Query Rewrite	297
16.1 Constraints and their impact on query optimization	297
16.2 Semantic Query Rewrite	297
16.3 Exploiting Uniqueness in Query Optimization	298
16.4 Bibliography	298
IV Plan Generation	299
17 Current Search Space and Its Limits	301
17.1 Plans with Outer Joins, Semijoins and Antijoins	301
17.2 Expensive Predicates and Functions	301
17.3 Techniques to Reduce the Search Space	301
17.4 Bibliography	301
18 Optimizing Queries with Disjunctions	303
18.1 Introduction	303
18.2 Using Disjunctive or Conjunctive Normal Forms	304
18.3 Bypass Plans	304
18.4 Implementation remarks	306
18.5 Other plan generators/query optimizer	306
18.6 Bibliography	307
19 Grouping and Aggregation	309
19.1 Introduction	309
19.2 Aggregate Functions	310
19.3 Normalization and Translation	313
19.3.1 Grouping and Aggregation in Different Query Languages	313
19.3.2 Extensions to Algebra	314
19.3.3 Normalization	314
19.3.4 Translation	314
19.4 Lazy and eager group by	314
19.5 Coalescing Grouping	316
19.6 More Possibilities	319
19.6.1 Eager/Lazy Group-By-Count	321
19.7 Translation into Our Notation	326
19.8 Aggregation of Non-Equi-Join Results	327
19.9 Bibliography	327
20 Grouping and Aggregation	329
20.1 Introduction	329
20.2 Lazy and eager group by	330
20.3 Coalescing Grouping	333
20.4 ToDo	336
21 Generating Plans for the Full Algebra	339
22 Generating DAG-structured Plans	341

23 Deriving and Dealing with Interesting Orderings and Groupings	343
23.1 Introduction	343
23.2 Problem Definition	344
23.2.1 Ordering	344
23.2.2 Grouping	346
23.2.3 Functional Dependencies	347
23.2.4 Algebraic Operators	347
23.2.5 Plan Generation	348
23.3 Overview	349
23.4 Detailed Algorithm	352
23.4.1 Overview	352
23.4.2 Determining the Input	353
23.4.3 Constructing the NFSM	354
23.4.4 Constructing the DFSM	356
23.4.5 Precomputing Values	357
23.4.6 During Plan Generation	358
23.4.7 Reducing the Size of the NFSM	358
23.4.8 Complex Ordering Requirements	361
23.5 Experimental Results	363
23.6 Total Impact	363
23.7 Influence of Groupings	365
23.8 Annotated Bibliography	369
24 Cardinality and Cost Estimates	373
24.1 Introduction	373
24.2 A First Approach	376
24.2.1 Top-Most Cost Formula (Overall Costs)	376
24.2.2 Summation of Operator Costs	376
24.2.3 CPU Cost	377
24.2.4 Abbreviations	377
24.2.5 I/O Costs	377
24.2.6 Cardinality Estimates	379
24.3 A First Logical Profile and its Propagation	381
24.3.1 The Logical Profile	381
24.3.2 Assumptions	382
24.3.3 Propagation for Selections	383
24.3.4 Profile Propagation for Joins	387
24.3.5 Profile Propagation for Projections	389
24.3.6 Profile Propagation for Division	392
24.3.7 Remarks	394
24.4 Approximation of a Set of Values	394
24.4.1 Approximations and Error Metrics	394
24.4.2 Example Applications	396
24.5 Approximation with Linear Models	396
24.5.1 Linear Models	396
24.5.2 Example Applications	401

24.5.3	Linear Models Under l_2	405
24.5.4	Linear Models Under l_∞	410
24.5.5	Linear Models Under l_q	414
24.5.6	Non-Linear Models	420
V	Implementation	423
25	Architecture of a Query Compiler	425
25.1	Compilation process	425
25.2	Architecture	425
25.3	Control Blocks	425
25.4	Memory Management	427
25.5	Tracing and Plan Visualization	427
25.6	Driver	427
25.7	Bibliography	427
26	Internal Representations	431
26.1	Requirements	431
26.2	Algebraic Representations	431
26.2.1	Graph Representations	432
26.2.2	Query Graph	432
26.2.3	Operator Graph	432
26.3	Query Graph Model (QGM)	432
26.4	Classification of Predicates	432
26.5	Treatment of Distinct	432
26.6	Query Analysis and Materialization of Analysis Results	432
26.7	Query and Plan Properties	433
26.8	Conversion to the Internal Representation	435
26.8.1	Preprocessing	435
26.8.2	Translation into the Internal Representation	435
26.9	Bibliography	435
27	Details on the Phases of Query Compilation	437
27.1	Parsing	437
27.2	Semantic Analysis, Normalization, Factorization, Constant Folding, and Translation	437
27.3	Normalization	439
27.4	Factorization	439
27.5	Constant Folding	440
27.6	Semantic analysis	440
27.7	Translation	442
27.8	Rewrite I	447
27.9	Plan Generation	447
27.10	Rewrite II	447
27.11	Code generation	447
27.12	Bibliography	448

28 Hard-Wired Algorithms	449
28.1 Hard-wired Dynamic Programming	449
28.1.1 Introduction	449
28.1.2 A plan generator for bushy trees	453
28.1.3 A plan generator for bushy trees and expensive selections	453
28.1.4 A plan generator for bushy trees, expensive selections and functions	455
28.2 Bibliography	455
29 Rule-Based Algorithms	457
29.1 Rule-based Dynamic Programming	457
29.2 Rule-based Memoization	457
29.3 Bibliography	457
30 Example Query Compiler	459
30.1 Research Prototypes	459
30.1.1 AQUA and COLA	459
30.1.2 Black Dahlia II	459
30.1.3 Epoq	459
30.1.4 Ereq	461
30.1.5 Exodus/Volcano/Cascade	462
30.1.6 Freytags regelbasierte System R-Emulation	464
30.1.7 Genesis	465
30.1.8 GOMbgo	467
30.1.9 Gral	470
30.1.10 Lambda-DB	473
30.1.11 Lanzelotte in short	473
30.1.12 Opt++	474
30.1.13 Postgres	474
30.1.14 Sciore & Sieg	476
30.1.15 Secondo	476
30.1.16 Squirrel	476
30.1.17 System R and System R*	478
30.1.18 Starburst and DB2	478
30.1.19 Der Optimierer von Straube	481
30.1.20 Other Query Optimizer	482
30.2 Commercial Query Compiler	484
30.2.1 The DB 2 Query Compiler	484
30.2.2 The Oracle Query Compiler	484
30.2.3 The SQL Server Query Compiler	488
VI Selected Topics	489
31 Generating Plans for Top-N-Queries?	491
31.1 Motivation and Introduction	491
31.2 Optimizing for the First Tuple	491

31.3	Optimizing for the First N Tuples	491
32	Recursive Queries	493
33	Issues Introduced by OQL	495
33.1	Type-Based Rewriting and Pointer Chasing Elimination	495
33.2	Class Hierarchies	497
33.3	Cardinalities and Cost Functions	499
34	Issues Introduced by XPath	501
34.1	A Naive XPath-Interpreter and its Problems	501
34.2	Dynamic Programming and Memoization	501
34.3	Naive Translation of XPath to Algebra	501
34.4	Pushing Duplicate Elimination	501
34.5	Avoiding Duplicate Work	501
34.6	Avoiding Duplicate Generation	501
34.7	Index Usage and Materialized Views	501
34.8	Cardinalities and Costs	501
34.9	Bibliography	501
35	Issues Introduced by XQuery	503
35.1	Reordering in Ordered Context	503
35.2	Result Construction	503
35.3	Unnesting Nested XQueries	503
35.4	Cardinalities and Cost Functions	503
35.5	Bibliography	503
36	Outlook	505
A	Query Languages?	507
A.1	Designing a query language	507
A.2	SQL	507
A.3	OQL	507
A.4	XPath	507
A.5	XQuery	507
A.6	Datalog	507
B	Query Execution Engine (?)	509
C	Glossary of Rewrite and Optimization Techniques	511
D	Useful Formulas	517
	Bibliography	518
	Index	578
E	ToDo	579

List of Figures

1.1	DBMS architecture	6
1.2	Query interpreter	6
1.3	Simple query interpreter	7
1.4	Query compiler	7
1.5	Query compiler architecture	8
1.6	Execution plan	10
1.7	Potential and actual search space	12
1.8	Generation vs. transformation	13
2.1	Relational algebra	17
2.2	Equivalences for the relational algebra	18
2.3	(Simplified) Canonical translation of SQL to algebra	19
2.4	Text book query optimization	20
2.5	Logical query optimization	21
2.6	Different join trees	27
2.7	Plans for example query (Part I)	28
2.8	Plans for example query (Part II)	29
2.9	Physical query optimization	30
2.10	Plan for example query after physical query optimization	30
3.1	Query graph for example query of Section 2.1	32
3.2	Query graph shapes	33
3.3	Illustrations for the IKKBZ Algorithm	55
3.4	A query graph, its directed join graph, some spanning trees and join trees	56
3.5	A query graph, its directed join tree, a spanning tree and its problem	58
3.6	Search space with sharing under optimality principle	63
3.7	Algorithm DPsize	70
3.8	Algorithm DPsub	72
3.9	Size of the search space for different graph structures	74
3.10	Algorithm DPccp	75
3.11	Enumeration Example for DPccp	75
3.12	Sample graph to illustrate EnumerateCsgRec	76
3.13	Call sequence for Figure 3.12	77
3.14	Example of rule transformations (RS-1)	99

3.15	Encoding Trees	104
3.16	Paths	105
3.17	Tree-merge	108
3.18	Algorithm <code>UnrankDecomposition</code>	110
3.19	Leaf-insertion	110
3.20	A query graph, its tree, and its standard decomposition graph . .	111
3.21	Algorithm <code>Adorn</code>	114
3.22	A query graph, a join tree, and its encoding	121
3.23	Pseudo code for <code>IDP-1</code>	124
3.24	Pseudocode for <code>IDP-2</code>	125
3.25	Subroutine <code>applicable-predicates</code>	127
3.26	Subroutine <code>construct-bushy-tree</code>	128
3.27	Subroutine <code>extract-plan</code> and its subroutine	128
3.28	Impact of selectivity on the search space	134
3.29	Impact of relation sizes on the search space	134
3.30	Impact of parameters on the performance of heuristics	134
3.31	Impact of selectivities on probabilistic procedures	135
4.1	Disk drive assembly	138
4.2	Disk drive read request processing	139
4.3	Time to read 100 MB from disk (depending on the number of 8 KB blocks read at once)	142
4.4	Time needed to read n random pages	144
4.5	Break-even point in fraction of total pages depending on page size	145
4.6	Physical organization of a relational database	148
4.7	Slotted pages and TIDs	150
4.8	Various physical record layouts	151
4.9	Clustered vs. non-clustered index	157
4.10	Illustration of seek cost estimate	190
5.1	Truth tables for two-valued logic	199
5.2	Simplification Rules	200
5.3	Laws for two-valued logic	201
5.4	Comparison functions in the presence of NULL values	201
5.5	Truth tables for three-valued logic	202
5.6	Interpretation of \perp values	202
7.1	Laws for Set Operations	205
7.2	Laws for Bag Operations	207
7.3	Outer join examples	221
7.4	Two equivalent plans	228
7.5	Applications of coalescing grouping	234
7.6	Example for Map Operator	236
7.7	Examples for Unary and Binary $\hat{\Gamma}$	237
11.1	Simplification rules for boolean expressions	258
11.2	Axioms for equality	258
11.3	265

11.4	Axioms for inequality	266
18.1	DNF plans	304
18.2	CNF plans	305
18.3	Bypass plans	305
19.1	Two equivalent plans	311
19.2	Applications of coalescing grouping	320
20.1	Two equivalent plans	331
20.2	Applications of coalescing grouping	337
23.1	Propagation of orderings and groupings	348
23.2	Possible FSM for orderings	350
23.3	Possible FSM for groupings	351
23.4	Combined FSM for orderings and groupings	351
23.5	Possible DFSM for Figure 23.4	351
23.6	Preparation steps of the algorithm	353
23.7	Initial NFSM for sample query	355
23.8	NFSM after adding D_{FD} edges	356
23.9	NFSM after pruning artificial states	356
23.10	Final NFSM	356
23.11	Resulting DFSM	357
23.12	<i>contains</i> Matrix	357
23.13	<i>transition</i> Matrix	357
23.14	Plan generation for different join graphs, Simmen's algorithm (left) vs. our algorithm (middle)	363
23.15	Memory consumption in KB for Figure 23.14	365
23.16	Time requirements for the preparation step	368
23.17	Space requirements for the preparation step	369
24.1	Overview of operations for cardinality and cost estimations	374
24.2	Calculating the lower bound D_G^\perp	391
24.3	Calculating the estimate for D_G	393
24.4	Q-error and plan optimality	405
24.5	Algorithm for best linear approximation under l_∞	414
24.6	Algorithm finding best linear approximation under l_q	419
25.1	The compilation process	426
25.2	Class Architecture of the Query Compiler	428
25.3	Control Block Structure	429
27.1	Expression	438
27.2	Expression hierarchy	439
27.3	Expression	440
27.4	Query 1	441
27.5	Internal representation	443
27.6	An algebraic operator tree with a d-join	446

27.7 Algebra	447
28.1 A sample execution plan	450
28.2 Different join operator trees	451
28.3 Bottom up plan generation	453
28.4 A Dynamic Programming Optimization Algorithm	454
30.1 Beispiel einer Epoq-Architektur	460
30.2 Exodus Optimierer Generator	462
30.3 Organisation der Optimierung	465
30.4 Ablauf der Optimierung	468
30.5 Architektur von GOMrbo	469
30.6 a) Architektur des Gral-Optimierers; b) Operatorhierarchie nach Kosten	470
30.7 Die Squiralarchitektur	477
30.8 Starburst Optimierer	479
30.9 Der Optimierer von Straube	481
33.1 Algebraic representation of a query	495
33.2 A join replacing pointer chasing	497
33.3 A Sample Class Hierarchy	498
33.4 Implementation of Extents	499

Preface

Goals

Primary Goals:

- book covers many query languages (at least SQL, OQL, XQuery (XPath))
- techniques should be represented as query language independent as possible
- book covers all stages of the query compilation process
- book completely covers fundamental issues
- book gives implementation details and tricks

Secondary Goals:

- book is thin
- book is not totally unreadable
- book separates concepts from implementation techniques

Organizing the material is not easy: The same topic pops up

- at different stages of the query compilation process and
- at different query languages

Acknowledgements

Introducer to query optimization: Günther von Bültzingsloewen

Peter Lockemann

First paper coauthor: Stefan Karl,

Coworkers: Alfons Kemper, Klaus Peithner, Michael Steinbrunn, Donald

Kossmann, Carsten Gerlhof, Jens Claussen,

Sophie Cluet, Vassilis Christophides, Georg Gottlob, V.S. Subramanian,

Sven Helmer, Birgitta König-Ries, Wolfgang Scheufele, Carl-Christian Kanne,

Thomas Neumann, Norman May, Matthias Brantner

Robin Aly

Discussions: Umesh Dayal, Dave Maier, Gail Mitchell, Stan Zdonik, Tamer Özsü, Arne Rosenthal,

Don Chamberlin, Bruce Lindsay, Guy Lohman, Mike Carey, Bennet Vance, Laura Haas, Mohan, CM Park,

Yannis Ioannidis, Götz Graefe, Serge Abiteboul, Claude Delobel Patrick Valduriez, Dana Florescu, Jerome Simeon, Mary Fernandez, Christoph Koch, Adam Bosworth, Joe Hellerstein, Paul Larson, Hennie Steenhagen, Harald Schöning, Bernhard Seeger,

Encouragement: Anand Deshpande

Manuscript: Simone Seeger,

and many others to be inserted.

Part I
Basics

Chapter 1

Introduction

1.1 General Remarks

Query languages like SQL or OQL are declarative. That is, they specify what the user wants to retrieve and not how to retrieve it. It is the task of the query compiler to generate a *query evaluation plan* (*evaluation plan* for short, or *execution plan* or simply *plan*) for a given query. The query evaluation plan (QEP) essentially is an operator tree with physical algebraic operators as nodes. It is evaluated by the runtime system. Figure 1.6 shows a detailed execution plan ready to be interpreted by the runtime system. Figure 28.1 shows an abstraction of a query plan often used to explain algorithms or optimization techniques.

The book tries to demystify query optimization and query optimizers. By means of the multi-lingual query optimizer BD II, the most important aspects of query optimizers and their implementation are discussed. We concentrate not only on the *query optimizer* core (Rewrite I, Plan Generator, Rewrite II) of the query compilation process but touch on all issues from parsing to code generation and quality assurance.

We start by giving a two-module overview of a database management system. One of these modules covers the functionality of the query compiler. The query compiler itself involves several submodules. For each submodule we discuss the features relevant for query compilation.

1.2 DBMS Architecture

Any database management system (DBMS) can be divided into two major parts: the compile time system (CTS) and the runtime system (RTS). User commands enter the compile time system which translates them into executables which are then interpreted by the runtime system (Fig. 1.1).

The input to the CTS are statements of several kinds, for example connect to a database (starts a session), disconnect from a database, create a database, drop a database, add/drop a schema, perform schema changes (add relations, object types, constraints, ...), add/drop indexes, run statistics commands, manually modify the DBMS statistics, begin of a transaction, end of a transac-

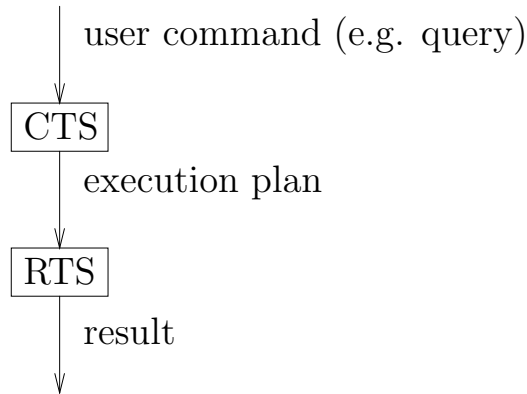


Figure 1.1: DBMS architecture




Figure 1.2: Query interpreter

tion, add/drop a view, update database items (e.g. tuples, relations, objects), change authorizations, and state a query. Within the book, we will only be concerned with the tiny last item.

1.3 Interpretation versus Compilation

There are two essential approaches to process a query: *interpretation* and *compilation*.

The path of a query through a query interpreter is illustrated in Figure 1.2. Query interpretation translates the query string into some internal representation that is mostly calculus-based. Optionally, some rewrite on this representation takes place. Typical steps during this rewrite phase are unnesting nested queries, pushing selections down, and introducing index structures. After that, the query is interpreted. A simple query interpreter is sketched in Figure 1.3. The first function, `interpret`, takes a simple SQL block and extracts the different clauses, initializes the result `R` and calls `eval`. Then, `eval` recursively evaluates the query by first producing the cross product of the entries in the `from` clause. After all of them have been processed, the predicate is applied and for those tuples where the `where` predicate evaluates to true, a result tuple is constructed and added to the result set `R`. Obviously, the sketched interpreter is far from being efficient. A much better approach has been described by Wong and Youssefi [808, 840].

Let us now discuss the compilation approach. The different steps are sum-

```

interprete(SQLBlock x) {

    /* possible rewrites go here */
    s := x.select();
    f := x.from();
    w := x.where();
    R :=  $\emptyset$ ; /* result */
    t := []; /* empty tuple */
    eval(s, f, w, t, R);
    return R;
}

eval(s, f, w, t, R) {

    if(f.empty())
        if(w(t))
            R += s(t);
    else
        foreach( $t' \in \text{first}(f)$ )
            eval(s, tail(f), w,  $t \circ t'$ , R);
}

```

Figure 1.3: Simple query interpreter




Figure 1.4: Query compiler

marized in Figure 1.4. First, the query is rewritten. Again, unnesting nested queries is a main technique for performance gains. Other rewrites will be discussed in Part ???. After the rewrite, the plan generation takes place. Here, an optimal plan is constructed. Whereas typically rewrite takes place on a calculus-based representation of the query, plan generation constructs an algebraic expression containing well-known operators like selection and join. Sometimes, after plan generation, the generated plan is refined: some polishing takes place. Then, code is generated, that can be interpreted by the runtime system. More specifically, the query execution engine—a part of the runtime system—interpretes the query execution plan. Let us illustrate this. The following query

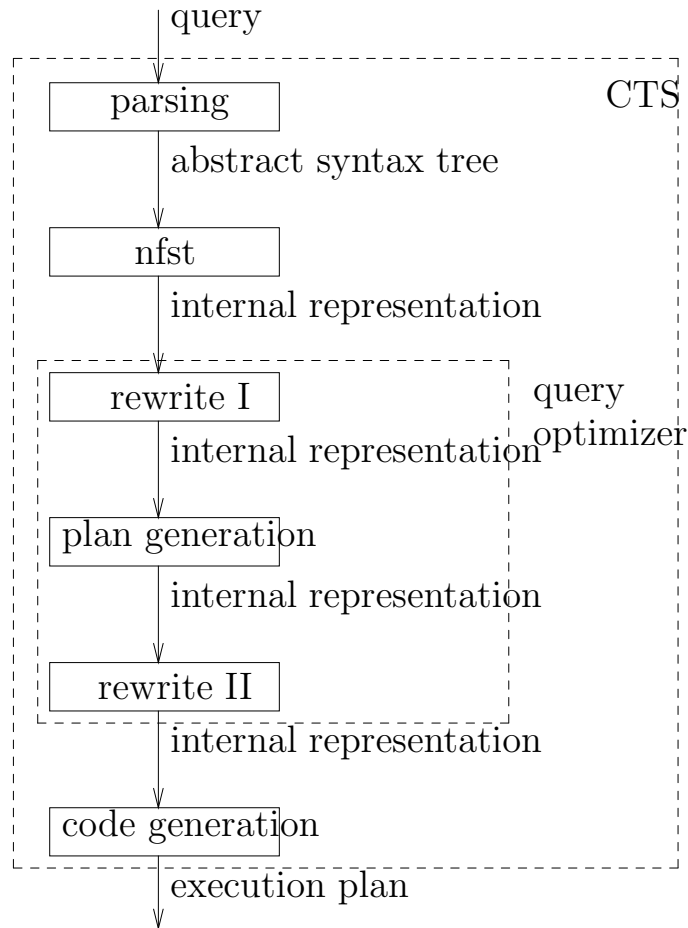


Figure 1.5: Query compiler architecture

is Query 1 of the now obsolete TPC-D benchmark [761].

```

SELECT   RETURNFLAG, LINESTATUS,
           SUM(QUANTITY) as SUM_QTY,
           SUM(EXTENDEDPRICE) as SUM_EXTPR,
           SUM(EXTENDEDPRICE * (1 - DISCOUNT)),
           SUM(EXTENDEDPRICE * (1 - DISCOUNT)*
              (1 + TAX)),
           AVG(QUANTITY),
           AVG(EXTENDEDPRICE),
           AVG(DISCOUNT),
           COUNT(*)
FROM     LINEITEM
WHERE    SHIPDDATE <= DATE '1998-12-01'
GROUP BY RETURNFLAG, LINESTATUS
ORDER BY RETURNFLAG, LINESTATUS
  
```

The CTS translates this query into a query execution plan. Part of the plan is shown in Fig. 1.6. One rarely sees a query execution plan. This is the reason why I included one. But note that the form of query execution plans differs from DBMS to DBMS since it is (unfortunately) not standardized the way SQL is. Most DBMSs can give the user abstract representations of query plans. It is worth the time to look at the plans generated by some commercial DBMSs.

I do not expect the reader to understand the plan in all details. Some of these details will become clear later. Anyway, this plan is given to the RTS which then interprets it. Part of the result of the interpretation might look like this:

RETURNFLAG	LINESTATUS	SUM_QTY	SUM_EXTPR	...
A	F	3773034	5319329289.68	...
N	F	100245	141459686.10	...
N	O	7464940	10518546073.98	...
R	F	3779140	5328886172.99	...

This should look familiar to you.

The above query plan is very simple. It contains only a few algebraic operators. Usually, more algebraic operators are present and the plan is given in a more abstract form that cannot be directly executed by the runtime system. Fig. 2.10 gives an example of an abstracted more complex operator tree. We will work with representations closer to this one.

A typical query compiler architecture is shown in Figure 1.5. The first component is the parser. It produces an abstract syntax tree. This is not always the case but this intermediate representation very much simplifies the task of following component. The NFST component performs several tasks. The first step is normalization. This mainly deals with introducing new variables for subexpressions. Factorization and semantic analysis are performed during NFST. Last, the abstract syntax tree is translated into the internal representation. All these steps can typically be performed during a single path through the query representation. Semantic analysis requires looking up schema definitions. This can be expensive and, hence, the number of lookups should be minimized. After NFST the core optimization steps rewrite I and plan generation take place. Rewrite II does some polishing before code generation. These modules directly correspond to the phases in Figure 1.4. They are typically further divided into submodules handling subphases. The most prominent example is the preparation phase that takes place just before the actual plan generation takes place. In our figures, we think of preparation as being part of the plan generation.

1.4 Requirements for a Query Compiler

Here are the main requirements for a query compiler:

1. Correctness
2. Completeness
3. Generate optimal plan (viz avoid the worst case)

```

(group
  (tbscan
    {segment 'lineitem.C4Kseg' 0 4096}
    {nalslottedpage 4096}
    {ctuple 'lineitem.cschema'}
    [ 20
      LOAD_PTR          1
      LOAD_SC1_C      8      1 2 // L_RETURNFLAG
      LOAD_SC1_C      9      1 3 // L_LINESTATUS
      LOAD_DAT_C     10      1 4 // L_SHIPDATE
      LEQ_DAT_ZC_C 4 '1998-02-09' 1
    ] 2 1 // number of help-registers and selection-register
  ) 10 22 // hash table size, number of registers
  [ // init
    MV_UI4_C_C      1          0 // COUNT(*) = 0
    LOAD_SF8_C      4          1 6 // L_QUANTITY
    LOAD_SF8_C      5          1 7 // L_EXTENDEDPRICE
    LOAD_SF8_C      6          1 8 // L_DISCOUNT
    LOAD_SF8_C      7          1 9 // L_TAX
    MV_SF8_Z_C      6          10 // SUM/AVG(L_QUANTITY)
    MV_SF8_Z_C      7          11 // SUM/AVG(L_EXTENDEDPRICE)
    MV_SF8_Z_C      8          12 // AVG(L_DISCOUNT)
    SUB_SF8_CZ_C    1.0        8 13 // 1 - L_DISCOUNT
    ADD_SF8_CZ_C    1.0        9 14 // 1 + L_TAX
    MUL_SF8_ZZ_C    7          13 15 // SUM(L_EXTDPRICE * (1 - L_DISC))
    MUL_SF8_ZZ_C    15         14 16 // SUM((...) * (1 + L_TAX))
  ] [ // advance
    INC_UI4         0          // inc COUNT(*)
    MV_PTR_Y        1          1
    LOAD_SF8_C      4          1 6 // L_QUANTITY
    LOAD_SF8_C      5          1 7 // L_EXTENDEDPRICE
    LOAD_SF8_C      6          1 8 // L_DISCOUNT
    LOAD_SF8_C      7          1 9 // L_TAX
    MV_SF8_Z_A      6          10 // SUM/AVG(L_QUANTITY)
    MV_SF8_Z_A      7          11 // SUM/AVG(L_EXTENDEDPRICE)
    MV_SF8_Z_A      8          12 // AVG(L_DISCOUNT)
    SUB_SF8_CZ_C    1.0        8 13 // 1 - L_DISCOUNT
    ADD_SF8_CZ_C    1.0        9 14 // 1 + L_TAX
    MUL_SF8_ZZ_B    7          13 17 15 // SUM(L_EXTDPRICE * (1 - L_DISC))
    MUL_SF8_ZZ_A    17         14 16 // SUM((...) * (1 + L_TAX))
  ] [ // finalize
    UIFC_C          0          18
    DIV_SF8_ZZ_C    10         18 19 // AVG(L_QUANTITY)
    DIV_SF8_ZZ_C    11         18 20 // AVG(L_EXTENDEDPRICE)
    DIV_SF8_ZZ_C    12         18 21 // AVG(L_DISCOUNT)
  ] [ // hash program
    HASH_SC1 2 HASH_SC1 3
  ] [ // compare program
    CMPA_SC1_ZY_C 2          2 0
    EXIT_NEQ      0
    CMPA_SC1_ZY_C 3          3 0
  ]
)

```

Figure 1.6: Execution plan

4. Efficiency, generate the plan fast, do not waste memory
5. Graceful degradation

6. Robustness

First of all, the query compiler must produce correct query evaluation plans. That is, the result of the query evaluation plan must be the result of the query as given by the specification of the query language. It must also cover the complete query language. The next issue is that an optimal query plan must (should) be generated. However, this is not always that easy. That is why some database researchers say that one must avoid the worst plan. Talking about the quality of a plan requires us to fix the optimization goal. Several goals are reasonable: We can optimize throughput, minimize response time, minimize resource consumption (both, memory and CPU), and so on. A good query compiler supports two optimization goals: minimize resource consumption and minimize the time to produce the first tuple. Obviously, both goals cannot be achieved at the same time. Hence, the query compiler must be instructed about the optimization goal.

Irrespective of the optimization goal, the query compiler should produce the query evaluation plan fast. It does not make sense to take 10 seconds to optimize a query whose execution time is below a second. This sounds reasonable but is not trivial to achieve. As we will see, the number of query execution plans that are equivalent to a given query, i.e. produce the same result as the query, can be very large. Sometimes, very large even means that not all plans can be considered. Taking the wrong approach to plan generation will result in no plan at all. This is the contrary of graceful degradation. Expressed positively, graceful degradation means that in case of limited resources, a plan is generated that may not be the optimal plan, but also not that far away from the optimal plan.

Last, typical software quality criteria should be met. We only mentioned robustness in our list, but others like maintainability must be met also.

1.5 Search Space

For a given query, there typically exists a high number of plans that are equivalent to the plan. Not all of these plans are accessible. Only those plans that can be generated by known optimization techniques (mainly algebraic equivalences) can potentially be generated. Since this number may still be too large, many query compilers restrict their search space further. We call the search space explored by a query optimizer the *actual search space*. The *potential search space* is the set of all plans that are known to be equivalent to the given query by applying the state of the art of query optimization techniques. The whole set of plans equivalent to a given query is typically unknown: we are not sure whether all optimization techniques have been discovered so far. Figure 1.7 illustrates the situation. Note that we run into problems if the actual search space is not a subset of the equivalent plans. Then the query compiler produces wrong results. As we will see in several chapters of this book, some optimization techniques have been proposed that produce plans that are not equivalent to the original query.

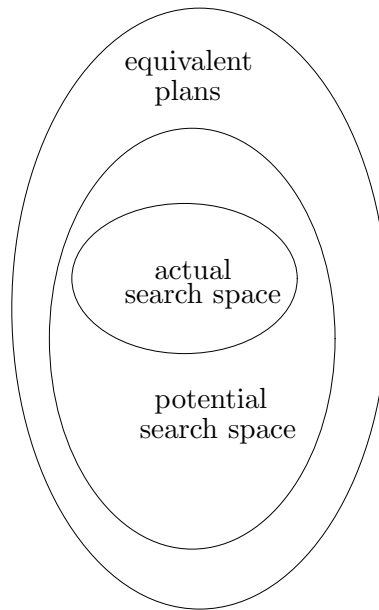


Figure 1.7: Potential and actual search space

1.6 Generation versus Transformation

Two different approaches to plan generation can be distinguished:

- The transformation-based approach transforms one query execution plan into another equivalent one. This can, for example, happen by applying an algebraic equivalence to a query execution plan in order to yield a better plan.
- The generic or synthetic approach produces a query execution plan by assembling building blocks and adding one algebraic operator after the other, until a complete query execution plan has been produced. Note that in this approach only when all building blocks and algebraic operators have been introduced the internal representation can be executed. Before that, no (complete) plan exists.

For an illustration see Figure 1.8.

A very important issue is how to explore the search space. Several well-known approaches exist: A*, Branch-and-bound, greedy algorithms, hill-climbing, dynamic programming, memoization, [179, 452, 453, 580]. These form the basis for most of the plan generation algorithms.

1.7 Focus

In this book, we consider only the compilation of queries. We leave out many special aspects like query optimization for multi-media database systems or

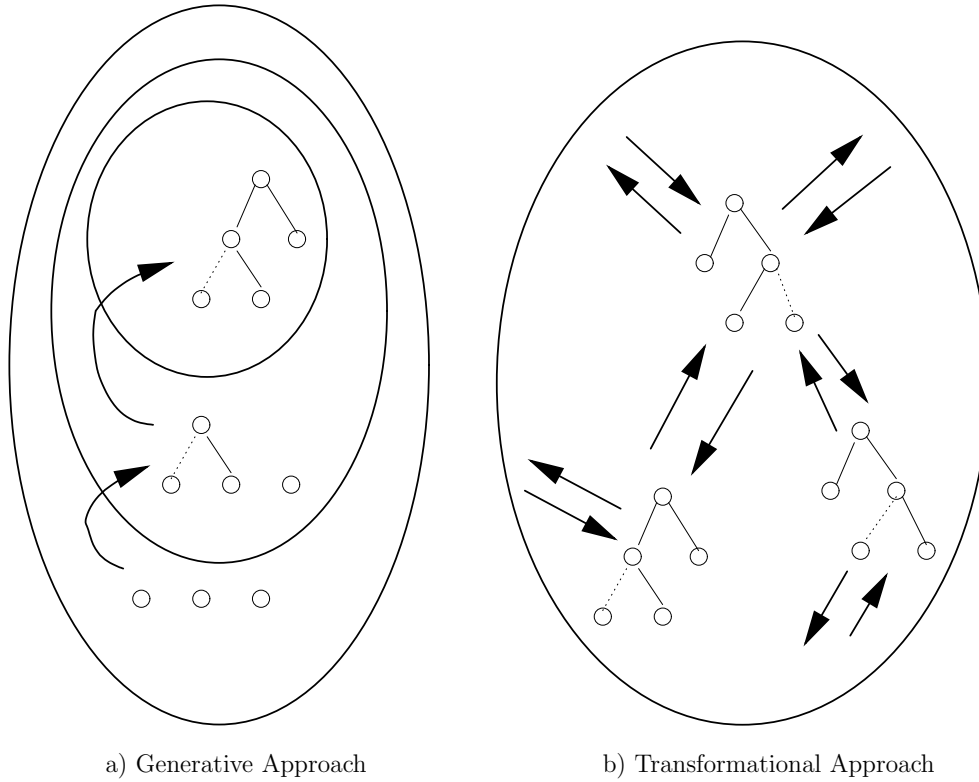


Figure 1.8: Generation vs. transformation

multidatabase systems. These are just two omissions. We further do not consider the translation of update statements which — especially in the presence of triggers — can become quite complex. Furthermore, we assume the reader to be familiar with the fundamentals of database systems [225, 413, 550, 604, 707] and their implementation [347, 270]. Especially, knowledge on query execution engines is required [294].

Last, the book presents a very personal view on query optimization. To see other views on the same topic, I strongly recommend to read the literature cited in this book and the references found therein. A good start are overview articles, PhD theses, and books, e.g. [782, 275, 381, 382, 398, 465] [522, 526, 562, 719, 739, 766, 767].

1.8 Organization of the Book

The first part of the book is an introduction to the topic. It should give an idea about the breadth and depth of query optimization. We first recapitulate query optimization the way it is described in numerous text books on database systems. There should be nothing really new here except for some pitfalls we will point out. The Chapter 3 is devoted to the join ordering problem. This has several reasons. First of all, at least one of the algorithms presented in this chapter forms the core of every plan generator. The second reason is that

this problem allows to discuss some issues like search space sizes and problem complexities. The third reason is that we do not have to delve into details. We can stick to very simple (you might call them unrealistic) cost functions, do not have to concern ourselves with details of the runtime system and the like. Expressed positively, we can concentrate on some algorithmic aspects of the problem. In Chapter 4 we do the opposite. The reader will not find any advanced algorithms in this chapter but plenty of details on disks and cost functions. The goal of the rest of the book is then to bring these issues together, broaden the scope of the chapters, and treat problems not even touched by them. The main issue not touched is query rewrite.

Chapter 2

Textbook Query Optimization

Almost every introductory textbook on database systems contains a section on query optimization (or at least query processing) [225, 413, 550, 604, 707]. Also, the two existing books on implementing database systems contain a section on query optimization [347, 270]. In this chapter we give an excerpt¹ of these sections and subsequently discuss the problems with the described approach. The bottom line will be that these descriptions of query optimization capture the essence of it but contain pitfalls that need to be pointed out and gaps to be filled.

2.1 Example Query and Outline

We use the following relations for our example query:

```
Student(SNo, SName, SAge, SYear)
Attend(ASNo, ALNo, AGrade)
Lecture(LNo, LTitle, LPNo)
Professor(PNo, PName)
```

Those attributes belonging to the key of the relations have been underlined.

With the following query we ask for all students attending a lecture by a Professor called “Larson”.

```
select distinct s.SName
from Student s, Attend a, Lecture l, Professor p
where s.SNo = a.ASNo and a.ALNo = l.LNo
      and l.LPNo = p.PNo and p.PName = ‘Larson’
```

The outline of the rest of the chapter is as follows. A query is typically translated into an algebraic expression. Hence, we first review the relational algebra and then discuss the translation process. Thereafter, we present the two phases of textbook query optimization: logical and physical query optimization. A brief discussion follows.

¹We do not claim to be fair to the above mentioned sections.

2.2 Algebra

Let us briefly recall the standard definition of the most important algebraic operators. Their inputs are relations, that is sets of tuples. Sets do not contain duplicates. The attributes of the tuples are assumed to be simple (non-decomposable) values. The most common algebraic operators are defined in Fig. 2.1. Although the common set operations union (\cup), intersection (\cap), and setdifference (\setminus) belong to the relational algebra, we did not list them. Remember that \cup and \cap are both commutative and associative. \setminus is neither of them. Further, for \cup and \cap , two distributivity laws hold. However, since these operations are not used in this section, we refer to Figure 7.1 in Section 7.1.1.

Before we can understand Figure 2.1, we must clarify some terms and notations. For us, a *tuple* is a mapping from a set of attribute names (or attributes for short) to their corresponding values. These values are taken from certain domains. An actual tuple is denoted embraced by brackets. They include a comma-separated list of the form attribute name, column and attribute value as in `[name: ‘Anton’, age: 2]`. If we have two tuples with different attribute names, they can be concatenated, i.e. we can take the union of their attributes. *Tuple concatenation* is denoted by ‘ \circ ’. For example `[name: ‘Anton’, age: 2] \circ [toy: ‘digger’]` results in `[name: ‘Anton’, age: 2, toy: ‘digger’]`. Let A and A' be two sets of attributes where $A' \subseteq A$ holds. Further let t a tuple with schema A . Then, we can project t on the attributes in A' (written as $t.A$). The resulting tuple contains only the attributes in A' ; others are discarded. For example, if t is the tuple `[name: ‘Anton’, age: 2, toy: ‘digger’]` and $A = \{\text{name, age}\}$, then $t.A$ is the tuple `[name: ‘Anton’, age: 2]`.

A relation is a set of tuples with the same attributes. The schema of a relation is the set of attributes. For a relation R this is sometimes denoted by $\text{sch}(R)$, the *schema* of R . We denote it by $\mathcal{A}(R)$ and extend it to any algebraic expression producing a set of tuples. That is, $\mathcal{A}(e)$ for any algebraic expression is the set of attributes the resulting relation defines. Consider the predicate `age = 2` where `age` is an attribute name. Then, `age` behaves like a free variable that must be bound to some value before the predicate can be evaluated. This motivates us to often use the terms *attribute* and *variable* synonymously. In the above predicate, we would call `age` a free variable. The set of free variables of an expression e is denoted by $\mathcal{F}(e)$.

Sometimes it is useful to work with sequences of attributes in comparison predicates. Let $A = \langle a_1, \dots, a_k \rangle$ and $B = \langle b_1, \dots, b_k \rangle$ be two attribute sequences. Then for any comparison operator $\theta \in \{=, \leq, <, \geq, >, \neq\}$, the expression $A\theta B$ abbreviates $a_1\theta b_1 \wedge a_2\theta b_2 \wedge \dots \wedge a_k\theta b_k$.

Often, a *natural join* is defined. Consider two relations R_1 and R_2 . Define $A_i := \mathcal{A}(R_i)$ for $i \in \{1, 2\}$, and $A := A_1 \cap A_2$. Assume that A is non-empty and $A = \langle a_1, \dots, a_n \rangle$. If A is non-empty, the natural join is defined as

$$R_1 \bowtie R_2 := \Pi_{A_1 \cup A_2}(R_1 \bowtie_p \rho_{A:A'}(R_2))$$

where $\rho_{A:A'}$ renames the attributes a_i in A to a'_i in A' and the predicate p has the form $A = A'$, i.e. $a_1 = a'_1 \wedge \dots \wedge a_n = a'_n$.

$$\begin{aligned}
\sigma_p(R) &:= \{r \mid r \in R, p(r)\} \\
\Pi_A(R) &:= \{r.A \mid r \in R\} \\
R_1 \times R_2 &:= \{r_1 \circ r_2 \mid r_1 \in R_1, r_2 \in R_2\} \\
R_1 \bowtie_p R_2 &:= \sigma_p(R_1 \times R_2)
\end{aligned}$$

Figure 2.1: Relational algebra

For our algebraic operators, several equivalences hold. They are given in Figure 2.2. For them to hold, we typically require that the relations involved have disjoint attribute sets. That is, we assume—even for the rest of the book—that attribute names are unique. This is often achieved by using the notation $R.a$ for a relation R or $v.a$ for a variable ranging over tuples with an attribute a . Another possibility is to use the renaming operator ρ .

Some equivalences are not always valid. Their validity depends on whether some condition(s) are satisfied or not. For example, Eqv. 2.4 requires $\mathcal{F}(p) \subseteq A$. That is, all attribute names occurring in p must be contained in the attribute set A the projection retains: otherwise, we could not evaluate p after the projection has been applied. Although all conditions in Fig. 2.2 are of this flavor, we will see throughout the course of the book that more complex conditions exist.

2.3 Canonical Translation

The next question is how to translate a given SQL query into the algebra. Since the relational algebra works on sets and not bags (multisets), we can only translate SQL queries that contain a **distinct**. Further, we restrict ourselves EX to flat queries not containing any subquery. Since negation, disjunction, aggregation, and quantifiers pose further problems, we neglect them. Further, we do not allow **group by**, **order by**, **union**, **intersection**, and **except** in our query. Last, we allow only attributes in the **select** clause and not more complex expressions.

Thus, the generic SQL query pattern we can translate into the algebra looks as follows:

```

select distinct  $a_1, a_2, \dots, a_m$ 
from  $R_1c_1, R_2c_2, \dots, R_nc_n$ 
where  $p$ 

```

Here, the R_i are relation names and the c_i are correlation names. The a_i in the **select** clause are attribute names (or expressions of the form $c_i.a_i$) taken from the relations in the **from** clause. The predicate p is assumed to be a conjunction of comparisons between attributes or attributes and constants.

The translation process then follows the procedure described in Figure 2.3. First, we construct an expression that produces the cross product of the entries

$$\begin{aligned} \sigma_{p_1 \wedge \dots \wedge p_k}(R) &\equiv \sigma_{p_1}(\dots(\sigma_{p_k}(R))\dots) & (2.1) \\ \sigma_{p_1}(\sigma_{p_2}(R)) &\equiv \sigma_{p_2}(\sigma_{p_1}(R)) & (2.2) \\ \Pi_{A_1}(\Pi_{A_2}(\dots(\Pi_{A_k}(R))\dots)) &\equiv \Pi_{A_1}(R) & (2.3) \\ &\text{if } A_i \subseteq A_j \text{ for } i < j \\ \Pi_A(\sigma_p(R)) &\equiv \sigma_p(\Pi_A(R)) & (2.4) \\ &\text{if } \mathcal{F}(p) \subseteq A \\ (R_1 \times R_2) \times R_3 &\equiv R_1 \times (R_2 \times R_3) & (2.5) \\ (R_1 \bowtie_{p_{1,2}} R_2) \bowtie_{p_{2,3}} R_3 &\equiv R_1 \bowtie_{p_{1,2}} (R_2 \bowtie_{p_{2,3}} R_3) & (2.6) \\ &\text{if } \mathcal{F}(p_{1,2}) \subseteq \mathcal{A}(R_1) \cup \mathcal{A}(R_2) \\ &\text{and } \mathcal{F}(p_{2,3}) \subseteq \mathcal{A}(R_2) \cup \mathcal{A}(R_3) \\ R_1 \times R_2 &\equiv R_1 \times R_2 & (2.7) \\ R_1 \bowtie_p R_2 &\equiv R_2 \bowtie_p R_1 & (2.8) \\ \sigma_p(R_1 \times R_2) &\equiv \sigma_p(R_1) \times R_2 & (2.9) \\ &\text{if } \mathcal{F}(p) \subseteq \mathcal{A}(R_1) \\ \sigma_p(R_1 \bowtie_q R_2) &\equiv \sigma_p(R_1) \bowtie_q R_2 & (2.10) \\ &\text{if } \mathcal{F}(p) \subseteq \mathcal{A}(R_1) \\ \Pi_A(R_1 \times R_2) &\equiv \Pi_{A_1}(R_1) \times \Pi_{A_2}(R_2) & (2.11) \\ &\text{if } A = A_1 \cup A_2, A_i \subseteq \mathcal{A}(R_i) \\ \Pi_A(R_1 \bowtie_p R_2) &\equiv \Pi_{A_1}(R_1) \bowtie_p \Pi_{A_2}(R_2) & (2.12) \\ &\text{if } \mathcal{F}(p) \subseteq A, A = A_1 \cup A_2, \\ &\text{and } A_i \subseteq \mathcal{A}(R_i) \\ \sigma_p(R_1 \theta R_2) &\equiv \sigma_p(R_1) \theta \sigma_p(R_2) & (2.13) \\ &\text{where } \theta \text{ is any of } \cup, \cap, \setminus \\ \Pi_A(R_1 \cup R_2) &\equiv \Pi_A(R_1) \cup \Pi_A(R_2) & (2.14) \\ \sigma_p(R_1 \times R_2) &\equiv R_1 \bowtie_p R_2 & (2.15) \end{aligned}$$

Figure 2.2: Equivalences for the relational algebra

found in the **from** clause. The result is

$$((\dots((R_1 \times R_2) \times R_3)\dots) \times R_n).$$

Next, we add a selection with the **where** predicate:

$$\sigma_p(((\dots((R_1 \times R_2) \times R_3)\dots) \times R_n)).$$

Last, we project on the attributes found in the **select** clause.

$$\Pi_{a_1, \dots, a_n}(\sigma_p(((\dots((R_1 \times R_2) \times R_3)\dots) \times R_n))).$$

For our example query

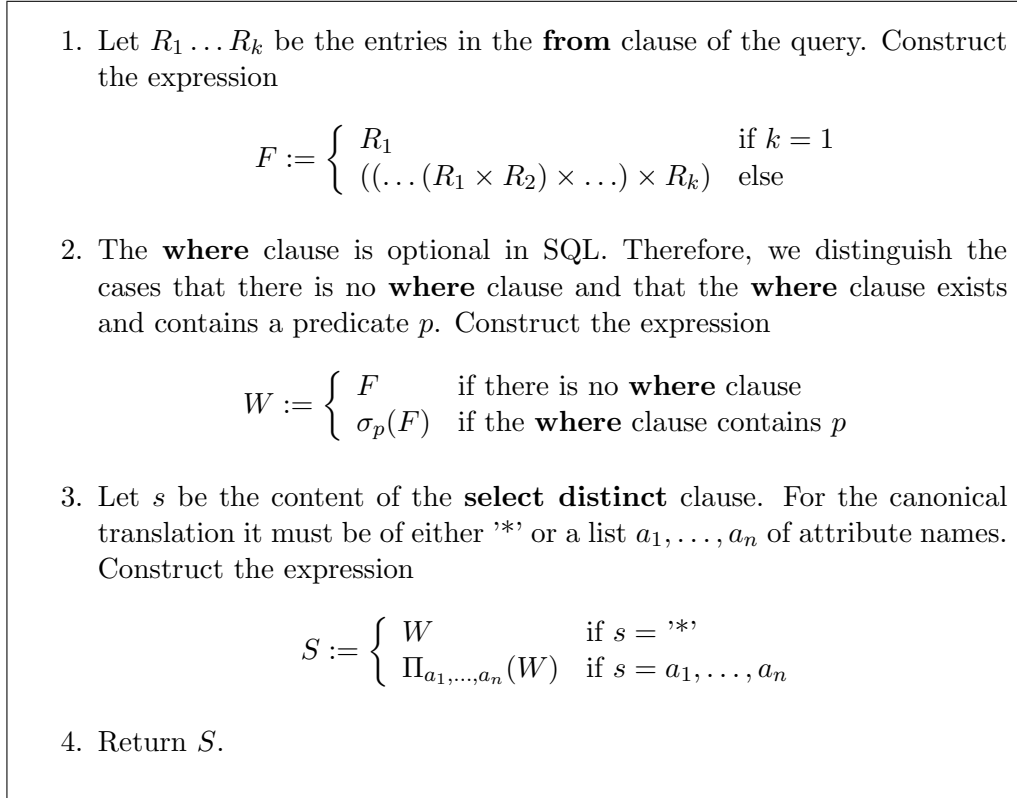


Figure 2.3: (Simplified) Canonical translation of SQL to algebra

```

select distinct s.SName
from Student s, Attend a, Lecture l, Professor p
where s.SNo = a.ASNo and a.ALNo = l.LNo
        and l.LPNo = p.PNo and p.PName = 'Larson'

```

the result of the translation is

$$\Pi_{s.SName}(\sigma_p(((Student[s] \times Attend[a]) \times Lecture[l]) \times Professor[p]))$$

where p equals

s.SNo = a.ASNo **and** a.ALNo = l.LNo **and** l.LPNo = p.PNo **and** p.PName = 'Larson'.

Note that we used the notation $R[r]$ to say that a relation named R has the correlation name r . During the course of the book we will be more precise about the semantics of this notation and it will deviate from the one suggested here. We will take r as a variable successively bound to the elements (tuples) in R . However, for the purpose of this chapter it is sufficient to think of it as associating a correlation name with a relation. The query is represented graphically in Figure 2.7 (top).

1. translate query into its canonical algebraic expression
2. perform logical query optimization
3. perform physical query optimization

Figure 2.4: Text book query optimization

2.4 Logical Query Optimization

Textbook query optimization takes place in two separate phases. The first phase is called *logical query optimization* and the second *physical query optimization*. Figure 2.4 lists all these steps together with the translation step. In this section we discuss logical query optimization. The foundation for this step is formed by the set of algebraic equivalences (see Figure 2.2). The set of algebraic equivalences spans the potential search space for this step. Given an initial algebraic expression—resulting from the translation of the given query—the algebraic equivalences can be used to derive all algebraic expressions that are equivalent to the initial algebraic expression. This set of all equivalent algebraic expressions can be derived by applying the equivalences first to the initial expression and then to all derived expressions until no new expression is derivable. Thereby, the algebraic equivalences can be applied in both directions: from left to right and from right to left. Care has to be taken that the conditions attached to the equivalences are obeyed.

Of course, whenever we find a new algebraic equivalence that could not be derived from those already known, adding this equivalence increases our potential search space. On the one hand, this has the advantage that in a larger search space we may find better plans. On the other hand, it increases the already large search space which might cause problems for its exploration. Nevertheless, finding new equivalences is a well-established sport among database researchers.

One remark on *better* plans. Plans can only be compared if costs can be attached to them via some cost function. This is what happens in most industrial strength query optimizers. However, at the level of logical algebraic expressions adding precise costs is not possible: too many implementation details are missing. These are added to the plan during the next phase called *physical query optimization*. As a consequence, we are left with plans without costs. The only thing we can do is to *heuristically* judge the effectiveness of applying an equivalence from left to right or in the opposite direction. As always with heuristics, the hope is that they work for most queries. However, it is typically very easy to find counter examples where the heuristics do not result in the best plan possible. (Again, *best* with respect to some metrics.) This finding can be generalized: any query optimization that takes place in more than a single phase risks missing the best plan. This is an important observation and we will come back to this issue more than once.

- | |
|---|
| <ol style="list-style-type: none"> 1. break up conjunctive selection predicates
(Eqv. 2.1: \rightarrow) 2. push down selections
(Eqv. 2.2: \rightarrow), (Eqv. 2.9: \rightarrow) 3. introduce joins
(Eqv. 2.15: \rightarrow) 4. determine join order
Eqv. 2.8, Eqv. 2.6, Eqv. 2.5, Eqv. 2.7 5. introduce and push down projections
(Eqv. 2.3: \leftarrow), (Eqv. 2.4: \rightarrow),
(Eqv. 2.11: \rightarrow), (Eqv. 2.12: \rightarrow) |
|---|

Figure 2.5: Logical query optimization

After these words of warning let us continue to discuss textbook query optimization. Logical query optimization requires the organization of all equivalences into groups. Further, the equivalences are directed. That is, it is fixed whether they are applied in a left to right or right to left manner. A *directed* equivalence is called *rewrite rule*. The groups of rewrite rules are then successively applied to the initial algebraic expression. Figure 2.5 describes the different steps performed during logical query optimization. Associated with each step is a set of rewrite rules that are applied to the input expression to yield a result expression. The numbers correspond to the equivalences in Figure 2.2. A small arrow indicates the direction in which the equivalences are applied.

The first step breaks up conjunctive selection predicates. The motivation behind this step is that selections with simple predicates can be moved around easier. The rewrite rule used in this step is Equivalence 2.1 applied from left to right. For our example query Step 1 results in

$$\Pi_{s.SName}(\sigma_{s.SNo=a.ASNo}(\sigma_{a.ALNo=l.LNo}(\sigma_{l.LPNo=p.PNo}(\sigma_{p.PName='Larson'}((Student[s] \times Attend[a]) \times Lecture[l]) \times Professor[p]))))))$$

The query is represented graphically in Figure 2.7 (middle).

Step 2 pushes selections down the operator tree. The motivation here is to reduce the number of tuples as early as possible such that subsequent (expensive) operators have smaller input. Applying this step to our example query yields:

$$\begin{aligned} & \Pi_{s.SName} (\\ & \quad \sigma_{l.LPNo=p.PNo} (\\ & \quad \quad \sigma_{a.ALNo=l.LNo} (\\ & \quad \quad \quad \sigma_{s.SNo=a.ASNo} (Student[s] \times Attend[a]) \\ & \quad \quad \quad \times Lecture[l]) \\ & \quad \quad \times (\sigma_{p.PName='Larson'} (Professor[p]))) \end{aligned}$$

The query is represented graphically in Figure 2.7 (bottom).

Excursion In general, we might encounter problems when pushing down selections. It may be the case that the order of the cross products is not well-suited for pushing selections down. If this is the case, we must consider reordering cross products during this step (Eqv. 2.7 and 2.5). To illustrate this point consider the following example query.

```
select distinct s.SName
from Student s, Lecture l, Attend a
where s.SNo = a.ASNo and a.ALNo = l.LNo
and l.LTitle = 'Databases I'
```

After translation and Steps 1 and 2 the algebraic expression looks like

$$\begin{aligned} & \Pi_{s.SName} (\\ & \quad \sigma_{s.SNo=a.ASNo} (\\ & \quad \quad \sigma_{a.ALNo=l.LNo} (\\ & \quad \quad \quad (Student[s] \times \sigma_{l.LTitle='Databases I'} (Lecture[l])) \times Attend[a]))). \end{aligned}$$

Neither of $\sigma_{s.SNo=a.ASNo}$ and $\sigma_{a.ALNo=l.LNo}$ can be pushed down further. Only after reordering the cross products such as in

$$\begin{aligned} & \Pi_{s.SName} (\\ & \quad \sigma_{s.SNo=a.ASNo} (\\ & \quad \quad \sigma_{a.ALNo=l.LNo} (\\ & \quad \quad \quad (Student[s] \times Attend[a]) \times \sigma_{l.LTitle='Databases I'} (Lecture[l]))) \end{aligned}$$

can $\sigma_{s.SNo=a.ASNo}$ be pushed down:

$$\begin{aligned} & \Pi_{s.SName} (\\ & \quad \sigma_{a.ALNo=l.LNo} (\\ & \quad \quad \sigma_{s.SNo=a.ASNo} (Student[s] \times Attend[a]) \\ & \quad \quad \times \sigma_{l.LTitle='Databases I'} (Lecture[l])) \end{aligned}$$

This is the reason why in some textbooks reorder cross products before selections are pushed down [225]. In this approach, reordering of cross products takes into account the selection predicates that can possibly be pushed down to the leaves and down to just prior a cross product. In any case, the Steps 2 and 4 are highly interdependent and there is no simple solution. \square

After this small excursion let us resume rewriting our main example query. The next step to be applied is converting cross products to join operations (Step 3). The motivation behind this step is that the evaluation of cross products

is very expensive and results in huge intermediate results. For every tuple in the left input an output tuple must be produced for every tuple in the right input. A join operation can be implemented much more efficiently. Applying Equivalence 2.15 from left to right to our example query results in

$$\begin{aligned} & \Pi_{s.SName} (\\ & \quad ((Student[s] \bowtie_{s.SNo=a.ASNo} Attend[a]) \\ & \quad \quad \bowtie_{a.ALNo=l.LNo} Lecture[l]) \\ & \quad \quad \quad \bowtie_{l.LPNo=p.PNo} (\sigma_{p.PName='Larson'}(Professor[p]))) \end{aligned}$$

The query is represented graphically in Figure 2.8 (top).

The next step is really tricky and involved: we have to find an optimal order for evaluating the joins. The join's associativity and commutativity gives us plenty of alternative (equivalent) evaluation plans. For our rather simple query Figure 2.6 lists some of the possible join orders where we left out the join predicates and used the single letter correlation names to denote the relations to be joined. Only p abbreviates the more complex expression $\sigma_{p.PName='Larson'}(Professor[p])$. The edges show how plans can be derived from other plans by applying commutativity (c) or associativity (a).

Unfortunately, we cannot ignore the problem of finding a good join order. It has been shown that the order in which joins are evaluated has an enormous influence on the total evaluation cost of a query. Thus, it is an important problem. On the other hand, the problem is really tough. Most join ordering problems turn out to be NP-hard. As a consequence, many different heuristics and cost-based algorithms have been invented. They are discussed in depth in Chapter 3. There we will also find examples showing how important (in terms of costs) the right choice of the join order is.

To continue with our example query, we use a very simple heuristics: among all possible joins select the one first that produces the smallest intermediate result. This can be motivated as follows. In our current algebraic expression, the first join to be executed is

$$Student[s] \bowtie_{s.SNo=a.ASNo} Attend[a].$$

All students and their attendances to some lecture are considered. The result and hence the input to the next join will be very big. On the other hand, if there is only one professor named *Larson*, the output of $\sigma_{p.PName='Larson'}(Professor[p])$ is a single tuple. Joining this single tuple with the relation **Lecture** results in an output containing one tuple for every lecture taught by *Larson*. For a large university, this will be a small subset of all lectures. Continuing this line, we get the following algebraic expression:

$$\begin{aligned} & \Pi_{s.SName} (\\ & \quad ((\sigma_{p.PName='Larson'}(Professor[p]) \\ & \quad \quad \bowtie_{p.PNo=l.LPNo} Lecture[l]) \\ & \quad \quad \quad \bowtie_{l.LNo=a.ALNo} Attend[a]) \\ & \quad \quad \quad \quad \bowtie_{a.ASNo=s.SNo} Student[s]) \end{aligned}$$

The query is represented graphically in Figure 2.8 (middle).

EX

The last step minimizes intermediate results by projecting out irrelevant attributes. An attribute is irrelevant, if it is not used further up the operator tree. When pushing down projections, we only apply them just before a pipeline breaker [294]. The reason is that for pipelined operators like selection, eliminating superfluous attributes does not gain much. The only pipeline breaker occurring in our plan is the join operator. Hence, before a join is applied, we project on the attributes that are further needed. The result is

$$\begin{aligned}
 & \Pi_{s.SName} (\\
 & \quad \Pi_{a.ASNo} (\\
 & \quad \quad \Pi_{l.LNo} (\\
 & \quad \quad \quad \Pi_{p.PNo} (\sigma_{p.PName='Larson'}(Professor[p])) \\
 & \quad \quad \quad \bowtie_{p.PNo=l.LPNo} \\
 & \quad \quad \quad \Pi_{l.LPNo,l.LNo}(Lecture[l])) \\
 & \quad \quad \quad \bowtie_{l.LNo=a.ALNo} \\
 & \quad \quad \quad \Pi_{a.ALNo,a.ASNo}(Attend[a])) \\
 & \quad \quad \quad \bowtie_{a.ASNo=s.SNo} \\
 & \quad \quad \quad \Pi_{s.SNo,s.SName}(Student[s]))
 \end{aligned}$$

This expression is represented graphically in Figure 2.8 (bottom).

2.5 Physical Query Optimization

Physical query optimization adds more information to the logical query evaluation plan. First, there exist many different ways to access the data stored in a database. One possibility is to scan a relation to find the relevant tuples. Another alternative is to use an index to access only the relevant parts. If an unclustered index is used, it might be beneficial to sort the *tuple identifiers* (TIDs²) to turn otherwise random disk accesses into sequential accesses. Since there is a multitude of possibilities to access data, this topic is discussed in depth in Chapter 4. Second, the algebraic operators used in the logical plan may have different alternative implementations. The most prominent example is the join operator that has many different implementations: simple nested loop join, blockwise nested loop join, blockwise nested loop join with in-memory hash table, index nested loop join, hybrid hash join, sort merge join, bandwidth join, special spatial joins, set joins, and structural joins. Most of these join implementations can be applied only in certain situations. Most algorithms only implement equi-joins where the join predicate is a conjunction of simple equalities. Further, all the implementations differ in cost and robustness. But also other operators like grouping may have alternative implementations. Typically, for these operators exist sort-based and hash-based alternatives. Third, some operators require certain *properties* for their input streams. For example, a sort merge join requires its input to be sorted on the join attributes occurring in the equalities of the join predicate. These attributes are called *join attributes*. The sortedness property can be *enforced* by a sort operator. The sort operator

²Sometimes TIDs are called RIDs (Row Identifiers).

is thus an *enforcer* since it makes sure that the required property holds. As we will see, properties and enforcers play a crucial role during plan generation.

If common subexpressions are detected at the algebraic level, it might be beneficial to compute them only once and store the result. To do so, a *tmp* operator must be introduced. Later on, we will see more of these operators that materialize (partial) intermediate results in order to avoid the same computation to be performed more than once. An alternative is to allow QEPs which are DAGs and not merely trees (see Section ??).

Physical query optimization is concerned with all the issues mentioned above. The outline of it is given in Figure 2.9. Let us demonstrate this for our small example query. Let us assume that there exists an index on the name of the professors. Then, instead of scanning the whole professor relation, it is beneficial to use the index to retrieve only those professors named *Larson*. Further, since a sort merge join is very robust and not the slowest alternative, we choose it as an implementation for all our join operations. This requires that we sort the inputs to the join operator on the join attributes. Since sorting is a pipeline breaker, we introduce it between the projections and the joins. The resulting plan is

$$\begin{aligned}
& \Pi_{s.SName}(\text{Sort}_{a.ASNo}(\Pi_{a.ASNo}(\text{Sort}_{l.LNo}(\Pi_{l.LNo}(\text{Sort}_{p.PNo}(\Pi_{p.PNo}(\text{IdxScan}_{p.PName='Larson'}(Professor[p]))) \\
& \quad \bowtie_{p.PNo=l.LPNo}^{smj} \\
& \quad \text{Sort}_{l.LPNo}(\Pi_{l.LPNo,l.LNo}(\text{Lecture}[l]))) \\
& \quad \bowtie_{l.LNo=a.ALNo}^{smj} \\
& \quad \text{Sort}_{a.ALNo}(\Pi_{a.ALNo,a.ASNo}(\text{Attend}[a]))) \\
& \quad \bowtie_{a.ASNo=s.SNo}^{smj} \\
& \text{Sort}_{s.SNo}(\Pi_{s.SNo,s.SName}(\text{Student}[s])))
\end{aligned}$$

where we annotated the joins with *smj* to indicate that they are sort merge joins. The *sort* operator has the attributes on which to sort as a subscript. We cheated a little bit with the notation of the index scan. The index is a physical entity stored in the database. An index scan typically allows to retrieve the TIDs of the tuples qualifying the predicate. If this is the case, another access to the relation itself is necessary to fetch the relevant attributes (*p.PNo* in our case) from the qualifying tuples of the relation. This issue is rectified in Chapter 4. The plan is shown as an operator graph in Figure 2.10.

2.6 Discussion

This chapter left open many interesting issues. We took it for granted that the presentation of a query is an algebraic expression or operator tree. Is this really true? We have been very vague about ordering joins and cross products. We only considered queries of the form **select distinct**. How can we assure correct duplicate treatment for **select all**? We separated query optimization into two distinct phases: logical and physical query optimization. Any separation into

different phases results in the danger of not producing an optimal plan. Logical query optimization turned out to be a little difficult: pushing selections down and reordering joins are mutually interdependent. How can we integrate these steps into a single one and thereby avoid the problem mentioned? Further, our logical query optimization was not cost based and cannot be: too much information is still missing from the plan to associate precise costs with a logical algebraic expression. How can we integrate the phases? How can we determine the costs of a plan? We covered only a small fraction of SQL. We did not discuss disjunction, negation, union, intersection, except, aggregate functions, group-by, order-by, quantifiers, outer joins, and nested queries. Furthermore, how about other query languages like OQL, XPath, XQuery? Further, enhancements like materialized views exist nowadays in many commercial systems. How can we exploit them beneficially? Can we exploit semantic information? Is our exploitation of index structures complete? What happens if we encounter NULL-values? Many questions and open issues remain. The rest of the book is about filling these gaps.




Figure 2.6: Different join trees

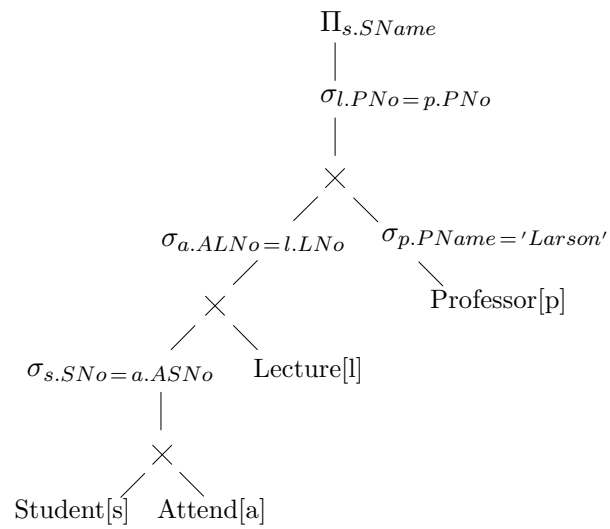
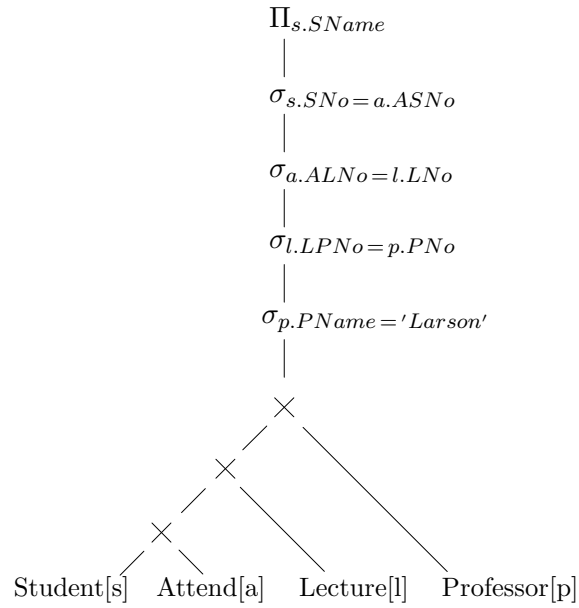
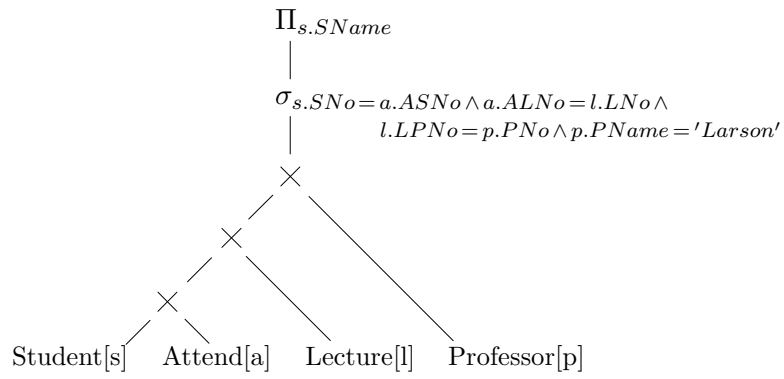


Figure 2.7: Plans for example query (Part I)

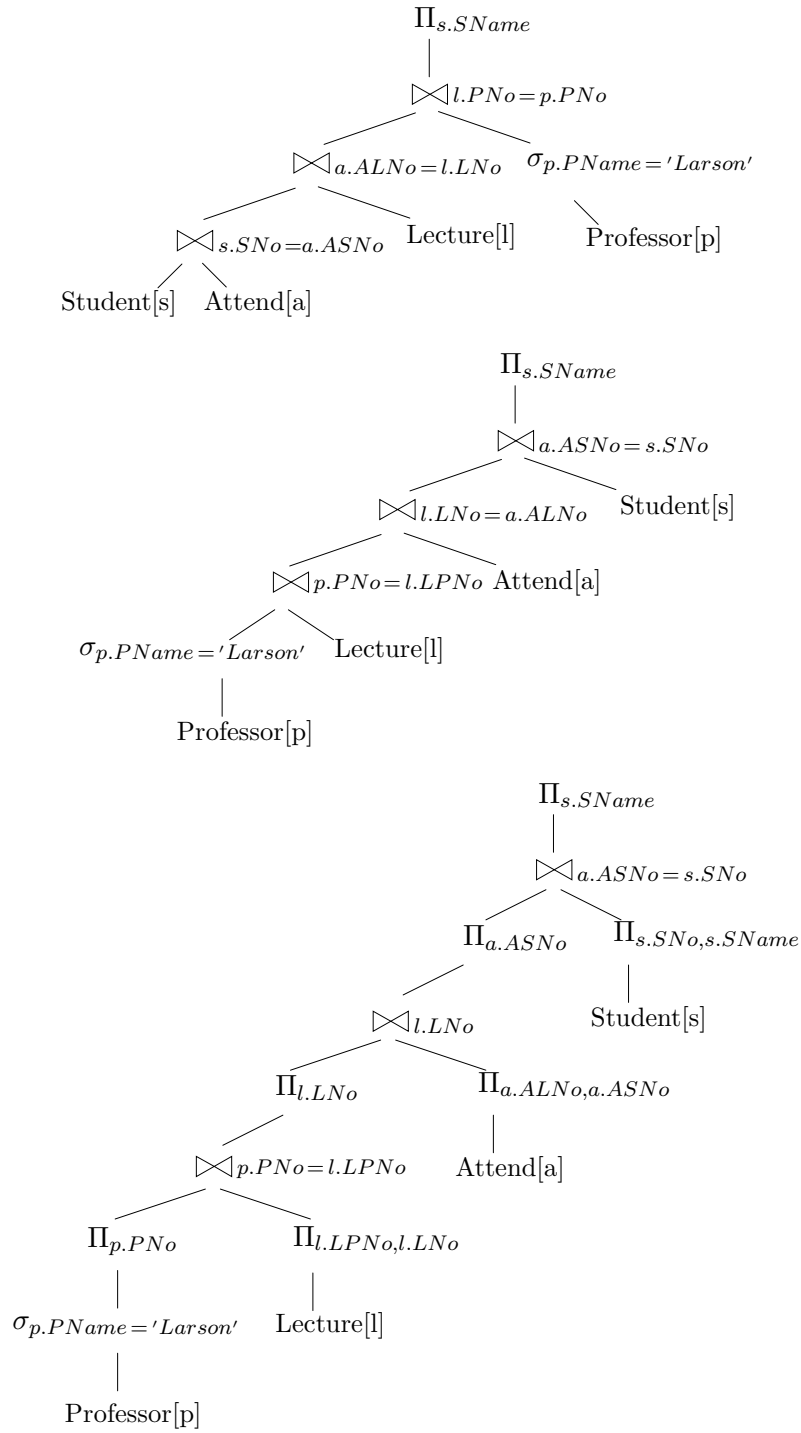


Figure 2.8: Plans for example query (Part II)

1. introduce index accesses
2. choose implementations for algebraic operators
3. introduce physical operators (sort, tmp)

Figure 2.9: Physical query optimization

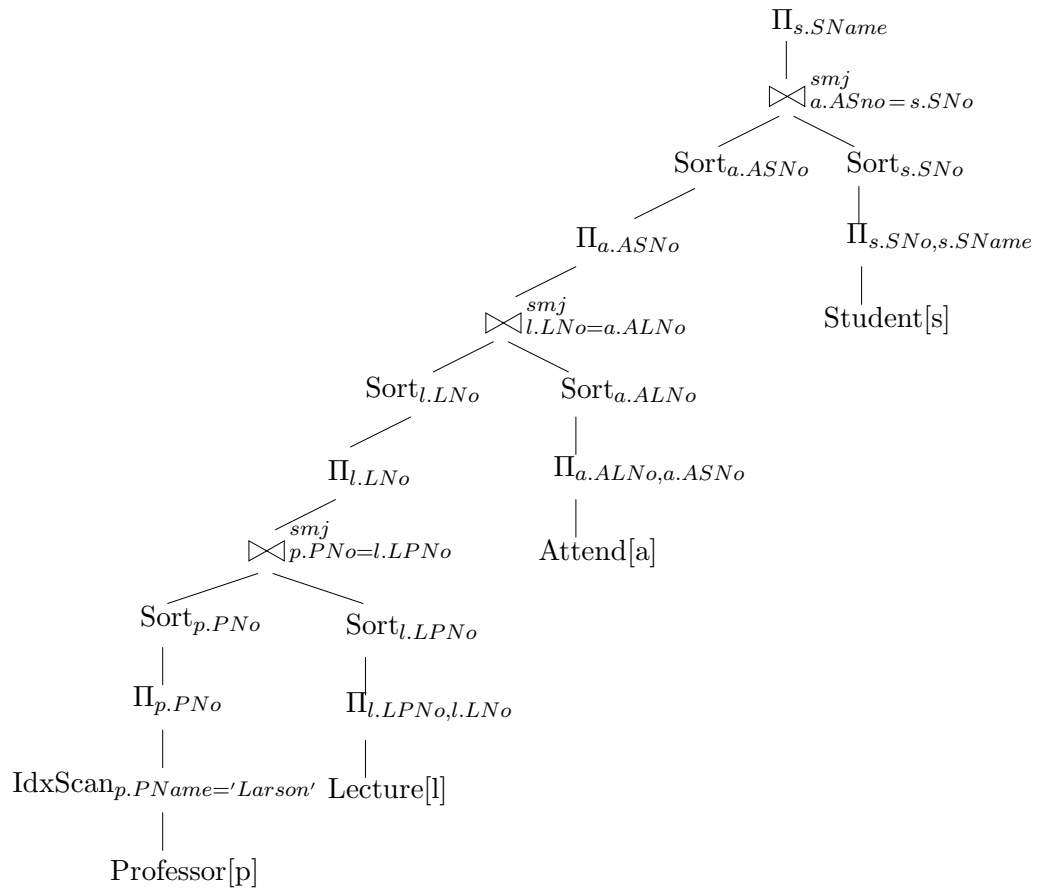


Figure 2.10: Plan for example query after physical query optimization

Chapter 3

Join Ordering

The problem of *join ordering* is a very restricted and — at the same time — a very complex one. We have touched this issue while discussing logical query optimization in Chapter 2. Join ordering is performed in Step 4 of Figure 2.5. In this chapter, we simplify the problem of join ordering by not considering duplicates, disjunctions, quantifiers, grouping, aggregation, or nested queries. Expressed positively, we concentrate on conjunctive queries with simple and cheap join predicates. What this exactly means will become clear in the next section. Subsequent sections discuss different algorithms for solving the join ordering problem. Finally, we take a look at the structure of the search space. This is important if different join ordering algorithms are compared via benchmarks. If the wrong parameters are chosen, benchmark results can be misleading.

The algorithms of this chapter form the core of every plan generator.

3.1 Queries Considered

A *conjunctive query* is one whose **where** clause contains a (complex) predicate which in turn is a conjunction of (simple) predicates. Hence, a conjunctive query involves only **and** and no *or* or *not* operations. A *simple predicate* is of the form $e_1\theta e_2$ where $\theta \in \{=, \neq, <, >, \leq, \geq\}$ is a comparison operator and the e_i are simple expressions in attribute names possibly containing some simple and cheap arithmetic operators. By cheap we mean that it is not worth applying extra optimization techniques. In this chapter, we restrict simple predicates even further to the form $A = B$ for attributes A and B . A and B must also belong to different relations such that every simple predicate in this chapter is a join predicate. There are two reasons for this restriction. First, the most efficient join algorithms rely on the fact that the join predicate is of the form $A = B$. Such joins are called *equi-joins*. Any other join is called a *non-equijoin*. Second, in relational systems joins on foreign key attributes of one relation and key attributes of the other relation are very common. Other joins are rare.

A *base relation* is a relation that is stored (explicitly) in the database. For the rest of the chapter, let R_i ($1 \leq i \leq n$) be n relations. These relations can be base relations but do not necessarily have to be. They could also be base relations to which predicates have already been supplied, e.g. as a result

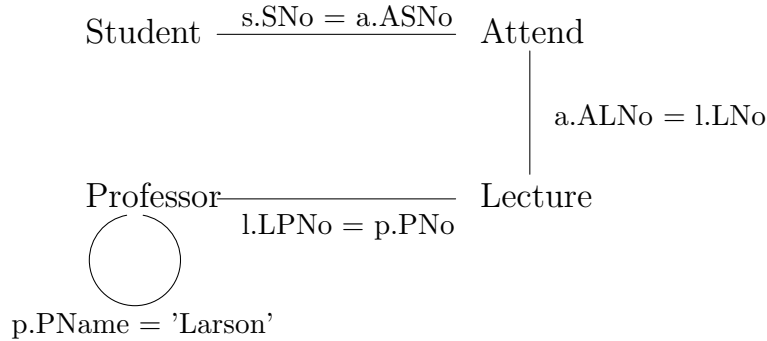


Figure 3.1: Query graph for example query of Section 2.1

of applying the first three steps of logical query optimization.

Summarizing, the queries we consider can be expressed in SQL as

```

select distinct *
from            $R_1, \dots, R_n$ 
where          $p$ 
  
```

where p is a conjunction of simple join predicates with attributes from exactly two relations. The latter restriction is not really necessary for the algorithms presented in this chapter but simplifies the exposition.

3.1.1 Query Graph

A query graph is a convenient representation of a query. It is an undirected graph with nodes R_1, \dots, R_n . For every simple predicate in the conjunction P whose attributes belong to the relations R_i and R_j , we add an edge between R_i and R_j . This edge is labeled by the simple predicate. From now on, we denote the join predicate connecting R_i and R_j by $p_{i,j}$. In general, $p_{i,j}$ can be a conjunction of simple join predicates connecting R_i and R_j .

If query graphs are used for more than join ordering, selections need to be represented. This is done by self-edges from the relation to which the selection applies to itself. For the example query of Chapter 2.6, Figure 3.1 contains the according query graph.

Query graphs can have many different shapes. The shapes that play a certain role in query optimization and the evaluation of join ordering algorithms are shown in Fig. 3.2. The query graph classes relevant for this chapter are chain queries, star queries, tree queries, cyclic queries and clique queries. Note that these classes are not disjoint and that some classes are subsets of other classes.

EX

In this chapter, we only treat connected query graphs. These can be evaluated without cross products.

Excursion In general, the query graph is a hypergraph [768] as the following example shows.

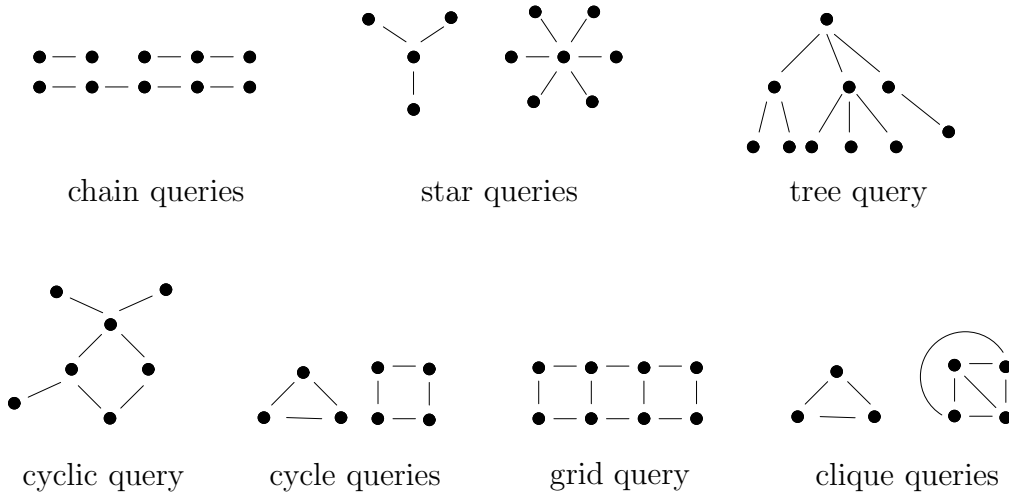


Figure 3.2: Query graph shapes

```

select *
from R1, R2, R3, R4
where f(R1.a, R2.a,R3.a) = g(R2.b,R3.b,R4.b)

```

3.1.2 Join Tree

A join tree is an algebraic expression in relation names and join operators. Sometimes, cross products are allowed, too. A cross product is the same as a join operator with *true* as its join predicate. A join tree has its name from its graph representation. There, a join tree is a binary tree whose leaf nodes are the relations and whose inner nodes are joins (and possibly cross products). The edges represent the input/output relationship. Examples of join trees have been shown in Figure 2.6.

Join trees fall into different classes. The most important classes are left-deep trees, right-deep trees, zig-zag trees, and bushy trees. *Left-deep trees* are join trees where every join has one of the relations R_i as its right input. *Right-deep trees* are defined analogously. In *zig-zag trees* at least one input of every join is a relation R_i . The class of zig-zag trees contains both left-deep and right-deep trees. For *bushy trees* no restriction applies. Hence, the class of bushy trees contains all of the above three classes. The roots of these notions date back to the paper by Selinger et al. [672], where the search space of the query optimizer was restricted to left-deep trees. There are two main reasons for this restriction. First, only one intermediate result is generated at any time during query evaluation. Second, the number of left-deep trees is far less than the number of e.g. bushy trees. The other classes were then added later by other researchers whenever they found better join trees in them. The different classes are illustrated in Figure 2.6. From left to right, the columns contain left-deep, zig-zag, right-deep, and bushy trees.

Left-deep trees directly correspond to an ordering (i.e. a permutation) of the relations. For example, the left-deep tree

$$((((R_2 \bowtie R_3) \bowtie R_1) \bowtie R_4) \bowtie R_5)$$

directly corresponds to the permutation R_2, R_3, R_1, R_4, R_5 . It should be clear that there is a one-to-one correspondence between permutations and left-deep join trees. We will also use the term *sequence of relations* synonymously. The notion of *join ordering* goes back to the times where only left-deep trees were considered and, hence, producing an optimal join tree was equivalent to optimally ordering the joins, i.e. determining a permutation with lowest cost.

Left-deep, right-deep, and zig-zag trees can be classed under the general term *linear trees*. Sometimes, the term linear trees is used synonymously for left-deep trees. We will not do so. Join trees are sometimes called *operator trees* or *query evaluation plans*. Although this is not totally wrong, these terms have a slightly different connotation. Operator trees typically contain more than only join operators. Query evaluation plans (QEPs or plans for short) typically have more information from physical query optimization associated with them.

3.1.3 Simple Cost Functions

In order to judge the quality of join trees, we need a cost function that associates a certain positive cost with each join tree. Then, the *task of join ordering* is to find among all equivalent join trees the join tree with lowest associated costs.

One part of any cost function are cardinality estimates. They are based on the cardinalities of the relations, i.e. the number of tuples contained in them. For a given relation R_i , we denote its cardinality by $|R_i|$.

Then, the cardinality of intermediate results must be estimated. This is done by introducing the notion of join selectivity. Let $p_{i,j}$ be a join predicate between relations R_i and R_j . The *selectivity* $f_{i,j}$ of $p_{i,j}$ is then defined as

$$f_{i,j} = \frac{|R_i \bowtie_{p_{i,j}} R_j|}{|R_i| * |R_j|}$$

This is the number of tuples in the join's result divided by the number of tuples in the Cartesian Product between R_i and R_j . If $f_{i,j}$ is 0.1, then only 10% of all tuples in the Cartesian Product survive the predicate $p_{i,j}$. Note that the selectivity is always a number between 0 and 1 and that $f_{i,j} = f_{j,i}$. We use an f and not an s , since the selectivity of a predicate is often called *filter factor*.

Besides the relation's cardinalities, the selectivities of the join predicates $p_{i,j}$ are assumed to be given as input to the join ordering algorithm. Therefore, we can compute the output cardinality of a join $R_i \bowtie_{p_{i,j}} R_j$, as

$$|R_i \bowtie_{p_{i,j}} R_j| = f_{i,j} |R_i| |R_j|$$

From this it becomes clear that if there is no join predicate for two relations R_i and R_j , we can assume a join predicate *true* and associate a selectivity of 1 with it. The output cardinality is then the cardinality of the cross product

between R_i and R_j . We also define $f_{i,i} = 1$ for all $1 \leq i \leq n$. This allows us to keep subsequent formulas simple.

We now need to extend our cardinality estimation to join trees. This can be done by recursively applying the above formula. Consider a join tree T joining two join trees T_1 and T_2 , i.e. $T = T_1 \bowtie T_2$. Then, the result cardinality $|T|$ can be calculated as follows. If T is a leaf R_i , then $|T| := |R_i|$. Otherwise,

$$|T| = \left(\prod_{R_i \in T_1, R_j \in T_2} f_{i,j} \right) |T_1| |T_2|.$$

Note that this formula assumes that the selectivities are independent of each other. Assuming independence is common but may be very misleading. More on this issue can be found in Chapter ???. Nevertheless, we assume independence and stick to the above formula.

For sequences of joins we can give a simple cardinality definition. Let $s = R_1, \dots, R_n$ be a sequence of relations. Then

$$|s| = \prod_{k=1}^n \left(\prod_{i=1}^k f_{i,k} |R_k| \right).$$

Given the above, a query graph alone is not really sufficient for the specification of a join ordering problem: cardinalities and selectivities are missing. On the other hand, from a complete list of cardinalities and selectivities we can derive the query graph. Obviously, the following defines a chain query with query graph $R_1 - - - R_2 - - - R_3$:

$$\begin{aligned} |R_1| &= 10 \\ |R_2| &= 100 \\ |R_3| &= 1000 \\ f_{1,2} &= 0.1 \\ f_{2,3} &= 0.2 \end{aligned}$$

In all examples, we assume for all i and j for which $f_{i,j}$ is not given that there is no join predicate and hence $f_{i,j} = 1$.

We now come to cost functions. The first cost function we consider is called C_{out} . For a join tree T , $C_{\text{out}}(T)$ is the sum of all output cardinalities of all joins in T . Recursively, we can define C_{out} as

$$C_{\text{out}}(T) = \begin{cases} 0 & \text{if } T \text{ is a single relation} \\ |T| + C_{\text{out}}(T_1) + C_{\text{out}}(T_2) & \text{if } T = T_1 \bowtie T_2 \end{cases}$$

From a theoretical point of view, C_{out} has many interesting properties: it is symmetric, it has the ASI property, and it can be applied to an expression of the logical algebra. From a practical point of view, however, it is rarely applied (yet).

In real cost functions, the cardinalities only serve as input to more complex formulas capturing the costs of a join implementation. Since real cost functions

are too complex for this section, we stick to simple cost functions proposed by Krishnamurthy, Boral, and Zaniolo [443]. They argue that these cost functions are appropriate for main memory database systems. For the three different join implementations nested loop join (nlj), hash join (hj), and sort merge join (smj), they give the following cost functions:

$$\begin{aligned} C_{\text{nlj}}(e_1 \bowtie_p e_2) &= |e_1||e_2| \\ C_{\text{hj}}(e_1 \bowtie_p e_2) &= h|e_1| \\ C_{\text{smj}}(e_1 \bowtie_p e_2) &= |e_1|\log(|e_1|) + |e_2|\log(|e_2|) \end{aligned}$$

where e_i are join trees and h is the average length of the collision chain in the hash table. We will assume $h = 1.2$. All these cost functions are defined for a single join operator. The cost of a join tree is defined as the sum of the costs of all joins it contains. We use the symbols C_x to also denote the costs of not only a single join but the costs of the whole tree. Hence, for sequences s of relations, we have

$$\begin{aligned} C_{\text{nlj}}(s) &= \sum_{i=2}^n |s_1, \dots, s_{i-1}| * |s_i| \\ C_{\text{hj}}(s) &= \sum_{i=2}^n 1.2|s_1, \dots, s_{i-1}| \\ C_{\text{smj}}(s) &= \sum_{i=2}^n |s_1, \dots, s_{i-1}| \log(|s_1, \dots, s_{i-1}|) + \sum_{i=1}^n |s_i| \log(|s_i|) \end{aligned}$$

Some notes on the cost functions are in order. First, note that these cost functions are even for main memory a little incomplete. For example, constant factors are missing. Second, the cost functions are mainly devised for left-deep trees. This becomes apparent when looking at the costs of hash joins. It is assumed that the right input is already stored in an appropriate hash table. Obviously, this can only hold for base relations, giving rise to left-deep trees. Third, C_{hj} and C_{smj} do not work for cross products. However, we can extend these cost functions by defining the cost of a cross product to be equal to its output cardinality, which happens to be the cost of C_{nlj} . Fourth, in reality, more complex cost functions are used and other parameters like the width of the tuples—i.e. the number of bytes needed to store them—also play an important role. Fifth, the above cost functions assume that the same join algorithm is chosen throughout the whole plan. In practice, this will not be true.

For the above chain query, we compute the costs of different join trees. The last join tree contains a cross product.

	C_{out}	C_{nlj}	C_{hj}	C_{smj}
$R_1 \bowtie R_2$	100	1000	12	697.61
$R_2 \bowtie R_3$	20000	100000	120	10630.26
$R_1 \times R_3$	10000	10000	10000	10000.00
$(R_1 \bowtie R_2) \bowtie R_3$	20100	101000	132	11327.86
$(R_2 \bowtie R_3) \bowtie R_1$	40000	300000	24120	32595.00
$(R_1 \times R_3) \bowtie R_2$	30000	1010000	22000	143542.00

For the calculation of C_{out} note that $|R_1 \bowtie R_2 \bowtie R_3| = 20000$ is included in all of the last three lines of its column. For the nested loop cost function, the costs are calculated as follows:

$$\begin{aligned} C_{\text{nlj}}((R_1 \bowtie R_2) \bowtie R_3) &= 1000 + 100 * 1000 = 101000 \\ C_{\text{nlj}}((R_2 \bowtie R_3) \bowtie R_1) &= 100000 + 20000 * 10 = 300000 \\ C_{\text{nlj}}((R_1 \times R_3) \bowtie R_2) &= 10000 + 10000 * 100 = 1010000 \end{aligned}$$

The reader should verify the other costs.

Several observations can be made from the above numbers:

- The costs of different join trees differ vastly under every cost function. Hence, it is worth spending some time to find a cheap join order.
- The costs of the same join tree differ under the different cost functions.
- The cheapest join tree is $(R_1 \bowtie R_2) \bowtie R_3$ under all four cost functions.
- Join trees with cross products are expensive. Thus, a heuristics often used is not to consider join trees that contain unnecessary cross products. (If the query graph consists of several unconnected components, then and only then cross products are necessary. In other words: if the query graph is connected, no cross products are necessary.)
- The join order matters even for join trees without cross products.

We would like to emphasize that the join order is also relevant under other cost functions.

Avoiding cross products is not always beneficial, as the following query specification shows:

$$\begin{aligned} |R_1| &= 1000 \\ |R_2| &= 2 \\ |R_3| &= 2 \\ f_{1,2} &= 0.1 \\ f_{1,3} &= 0.1 \end{aligned}$$

For C_{out} we have costs

Join Tree	C_{out}
$R_1 \bowtie R_2$	200
$R_2 \times R_3$	4
$R_1 \bowtie R_3$	200
$(R_1 \bowtie R_2) \bowtie R_3$	240
$(R_2 \times R_3) \bowtie R_1$	44
$(R_1 \bowtie R_3) \bowtie R_2$	240

Note that although the absolute numbers are quite small, the ratio of the best and the second best join tree is quite large. The reader is advised to find more examples and to apply other cost functions.

The following example illustrates that a bushy tree can be superior to any linear tree. Let us use the following query specification:

$$\begin{aligned}
 |R_1| &= 10 \\
 |R_2| &= 20 \\
 |R_3| &= 20 \\
 |R_4| &= 10 \\
 f_{1,2} &= 0.01 \\
 f_{2,3} &= 0.5 \\
 f_{3,4} &= 0.01
 \end{aligned}$$

If we do not consider cross products, we have for the symmetric (see below) cost function C_{out} the following join trees and costs:

Join Tree	C_{out}
$R_1 \bowtie R_2$	2
$R_2 \bowtie R_3$	200
$R_3 \bowtie R_4$	2
$((R_1 \bowtie R_2) \bowtie R_3) \bowtie R_4$	24
$((R_2 \bowtie R_3) \bowtie R_1) \bowtie R_4$	222
$(R_1 \bowtie R_2) \bowtie (R_3 \bowtie R_4)$	6

Note that all other linear join trees fall into one of these classes, due to the symmetry of the cost function and the join ordering problem. Again, the reader is advised to find more examples and to apply other cost functions.

If we want to annotate a join operator by its implementation—which is necessary for the correct computation of costs—we write \bowtie^{impl} for an implementation `impl`. For example, \bowtie^{smj} is a sort-merge join, and the according cost function C_{smj} is used to compute its costs.

Two properties of cost functions have some impact on the join ordering problem. The first is symmetry. A cost function C_{impl} is called *symmetric* if $C_{\text{impl}}(R_1 \bowtie^{\text{impl}} R_2) = C_{\text{impl}}(R_2 \bowtie^{\text{impl}} R_1)$ for all relations R_1 and R_2 . For symmetric cost functions, it does not make sense to consider commutativity. Hence, it suffices to consider left-deep trees only if we want to restrict ourselves to linear join trees. Note that C_{out} , C_{nlj} , C_{smj} are symmetric while C_{hj} is not.

The other property is the *adjacent sequence interchange* (ASI) property. Informally, the ASI property states that there exists a rank function such that the order of two subsequences is optimal if they are ordered according to the rank function. The ASI property is formally defined in Section 3.2.2. Only for tree queries and cost functions with the ASI property, a polynomial algorithm to find an optimal join order is known. Our cost functions C_{out} and C_{hj} have the ASI property, C_{smj} does not. Summarizing the properties of our cost functions, we see that the classification is orthogonal:

	ASI	\neg ASI
symmetric	$C_{\text{out}}, C_{\text{nlj}}$	C_{smj}
\neg symmetric	C_{hj}	(see text)

For the missing case of a non-symmetric cost function not having the ASI property, we can use the cost function of the hybrid hash join [204, 576].

We turn to another not really well-researched topic. The goal is to cut down the number of cost functions which have to be considered for optimization and to possibly allow for simpler cost functions, which saves time during plan generation. Unfortunately, we have to restrict ourselves to left-deep join trees. Let s denote a sequence or permutation of a given set of joins. We define an equivalence relation on cost functions.

Definition 3.1.1 *Let C and C' be two cost functions. Then*

$$C \equiv C' :\prec\rangle (\forall s \ C(s) \text{ minimal} \prec\rangle C'(s) \text{ minimal})$$

Here, s is a join sequence.

Obviously, \equiv is an equivalence relation.

Now we can define the ΣIR property.

Definition 3.1.2 *A cost function C is ΣIR : $\prec\rangle C \equiv C_{\text{out}}$.*

That is, ΣIR is the set of all cost functions that are equivalent to C_{out} .

Let us consider a very simple example. The last element of the sum in C_{out} is the size of the final join (all relations are joined). This is not the case for the following cost function:

$$C'_{\text{out}}(s) := \sum_{i=2}^{n-1} |s_1, \dots, s_i|$$

Obviously, we have C'_{out} is ΣIR . The next observation shows that we can construct quite complex ΣIR cost functions:

Observation 3.1.3 *Let C_1 and C_2 be two ΣIR cost functions. For non-decreasing functions $f_1 : R \rightarrow R$ and $f_2 : R \times R \rightarrow R$ and constants $c \in R$ and $d \in R^+$, we have that*

EX

$$\begin{aligned} &C_1 + c \\ &C_1 * d \\ &f_1 \circ C_1 \\ &f_2 \circ (C_1, C_2) \end{aligned}$$

are ΣIR . Here, \circ denotes function composition and (\cdot, \cdot) function pairing.

There are of course many more possibilities of constructing ΣIR functions. For the cost functions C_{hj} , C_{smj} , and C_{nlj} , we now investigate which of them have the ΣIR property.

Let us consider C_{hj} first. From

$$\begin{aligned} C_{hj}(s) &= \sum_{i=2}^n 1.2|s_1, \dots, s_{i-1}| \\ &= 1.2|s_1| + 1.2 \sum_{i=2}^{n-1} |s_1, \dots, s_i| \\ &= 1.2|s_1| + 1.2C'_{out}(s) \end{aligned}$$

and observation 3.1.3, we conclude that C_{hj} is ΣIR for a fixed relation to be joined first. If we can optimize C_{out} in polynomial time, then we can optimize C_{out} for a fixed starting relation. Indeed, by trying each relation as a starting relation, we can find the optimal join tree in polynomial time. An algorithm that computes the optimal solution for an arbitrary relation to be joined first can be found in Section 3.2.2.

EX

Now, consider C_{smj} . Since

$$\sum_{i=2}^n |s_1, \dots, s_{i-1}| \log(|s_1, \dots, s_{i-1}|)$$

is minimal if and only if

$$\sum_{i=2}^n |s_1, \dots, s_{i-1}|$$

is minimal and $\sum_{i=2}^n |s_i| \log(|s_i|)$ is independent of the order of the relations within s — that is constant — we conclude that C_{smj} is ΣIR .

Last, we have that C_{nlj} is not ΣIR . To see this, consider the following counter example with three relations R_1 , R_2 , and R_3 of sizes 10, 10, and 100, resp. The selectivities are $f_{1,2} = \frac{9}{10}$, $f_{2,3} = \frac{1}{10}$, and $f_{1,3} = \frac{1}{10}$. Now,

$$\begin{aligned} |R_1R_2| &= 90 \\ |R_1R_3| &= 100 \\ |R_2R_3| &= 100 \end{aligned}$$

and

$$\begin{aligned} C_{nl}(R_1R_2R_3) &= 10 * 10 + 90 * 100 = 9100 \\ C_{nl}(R_1R_3R_2) &= 10 * 100 + 100 * 10 = 2000 \\ C_{nl}(R_2R_3R_1) &= 10 * 100 + 100 * 10 = 2000 \end{aligned}$$

We see that $R_1R_2R_3$ has the smallest sum of intermediate result sizes but produces the highest cost. Hence, C_{nlj} is not ΣIR .

3.1.4 Classification of Join Ordering Problems

After having discussed the different classes of query graphs, join trees and cost functions, we can classify join ordering problems. To define a certain join ordering problem, we have to pick one entry from every class:

Query Graph Classes \times Possible Join Tree Classes \times Cost Function Classes

The query graph classes considered are *chain*, *star*, *tree*, and *cyclic*. For the join tree classes we distinguish between the different join tree shapes, i.e. whether they are left-deep, zig-zag, or bushy trees. We left out the right-deep trees, since they do not differ in their behavior from left-deep trees. We also have to take into account whether cross products are considered or not. For cost functions, we use a simple classification: we only distinguish between those that have the ASI property and those that do not. This leaves us with $4 \times 3 \times 2 \times 2 = 48$ different join ordering problems. For these, we will first review search space sizes and complexity. Then, we discuss several algorithms for join ordering. Last, we give some insight into cost distributions over the search space and how this might influence the benchmarking of different join ordering algorithms.

3.1.5 Search Space Sizes

Since search space sizes are easier to count if cross products are allowed, we consider them first. Then we turn to search spaces where cross products are not considered.

Join Trees with Cross Products We consider the number of join trees for a query graph with n relations. When cross products are allowed, the number of left-deep and right-deep join trees is $n!$. By allowing cross products, the query graph does not restrict the search space in any way. Hence, any of the $n!$ permutations of the n relations corresponds to a valid left-deep join tree. This is true independent of the query graph.

Similarly, the number of zig-zag trees can be estimated independently of the query graph. First note that for joining n relations, we need $n - 1$ join operators. From any left-deep tree, we derive zig-zag trees by using the join's commutativity and exchange the left and right inputs. Hence, from any left-deep tree for n relations, we can derive 2^{n-2} zig-zag trees. We have to subtract another one, since the bottommost joins' arguments are exchanged in different left-deep trees. Thus, there exists a total of $2^{n-2}n!$ zig-zag trees. Again, this number is independent of the query graph.

The number of bushy trees can be estimated as follows. First, we need the number of binary trees. For n leaf nodes, the number of binary trees is given by $\mathcal{C}(n - 1)$, where $\mathcal{C}(n)$ is defined by the recurrence

$$\mathcal{C}(n) = \sum_{k=0}^{n-1} \mathcal{C}(k)\mathcal{C}(n - k - 1)$$

with $\mathcal{C}(0) = 1$. The numbers $\mathcal{C}(n)$ are called the *Catalan Numbers* (see [179]). They can also be computed by the following formula:

$$\mathcal{C}(n) = \frac{1}{n + 1} \binom{2n}{n}.$$

The Catalan Numbers grow in the order of $\Theta(4^n/n^{3/2})$.

After we know the number of binary trees with n leaves, we now have to attach the n relations to the leaves in all possible ways. For a given binary tree, this can be done in $n!$ ways. Hence, the total number of bushy trees is $n!\mathcal{C}(n-1)$. This can be simplified as follows (see also [262, 454, 753]):

$$\begin{aligned} n!\mathcal{C}(n-1) &= n! \frac{1}{n} \binom{2(n-1)}{n-1} \\ &= n! \frac{1}{n} \frac{(2n-2)!}{(n-1)!((2n-2)-(n-1))!} \\ &= \frac{(2n-2)!}{(n-1)!} \end{aligned}$$

Chain Queries, Left-Deep Join Trees, No Cartesian Product We now derive the function that calculates the number of left-deep join trees with no cross products for a chain query of n relations. That is, the query graph is $R_1 - R_2 - \dots - R_{n-1} - R_n$. Let us denote the number of join trees by $f(n)$. Obviously, for $n = 0$ there is only one (the empty) join tree. For $n = 1$, there is also only one join tree (no join). For larger n : Consider the join trees for $R_1 - \dots - R_{n-1}$ where relation R_{n-1} is the k -th relation from the bottom where k ranges from 1 to $n-1$. From such a join tree we can derive join trees for all n relations by adding relation R_n at any position following R_{n-1} . There are $n-k$ such join trees. Only for $k = 1$, we can also add R_n below R_{n-1} . Hence, for $k = 1$ we have n join trees. How many join trees with R_{n-1} at position k are there? For $k = 1$, R_{n-1} must be the first relation to be joined. Since we do not consider cross products, it must be joined with R_{n-2} . The next relation must be R_{n-3} , and so on. Hence, there is only one such join tree. For $k = 2$, the first relation must be R_{n-2} , which is then joined with R_{n-1} . Then R_{n-3}, \dots, R_1 must follow in this order. Again, there is only one such join tree. For higher k , for R_{n-1} to occur safely at position k (no cross products) the $k-1$ relations R_{n-2}, \dots, R_{n-k} must occur before R_{n-1} . There are exactly $f(k-1)$ join trees for the $k-1$ relations. On each such join tree we just have to add R_{n-1} on top of it to yield a join tree with R_{n-1} at position k .

Now we can compute the $f(n)$ as $n + \sum_{k=2}^{n-1} f(k-1) * (n-k)$ for $n > 1$. Solving this recurrence gives us $f(n) = 2^{n-1}$. The proof is by induction. The case $n = 1$ is trivial.

The induction step for $n > 1$ provided by Thomas Neumann goes as follows:

$$\begin{aligned}
f(n) &= n + \sum_{k=2}^{n-1} f(k-1) * (n-k) \\
&= n + \sum_{k=0}^{n-3} f(k+1) * (n-k-2) \\
&= n + \sum_{k=0}^{n-3} 2^k * (n-k-2) \\
&= n + \sum_{k=1}^{n-2} k 2^{n-k-2} \\
&= n + \sum_{k=1}^{n-2} 2^{n-k-2} + \sum_{k=2}^{n-2} (k-1) 2^{n-k-2} \\
&= n + \sum_{i=1}^{n-2} \sum_{j=i}^{n-2} 2^{n-j-2} \\
&= n + \sum_{i=1}^{n-2} \sum_{j=0}^{n-i-2} 2^j \\
&= n + \sum_{i=1}^{n-2} (2^{n-i-1} - 1) \\
&= n + \sum_{i=1}^{n-2} 2^i - (n-2) \\
&= n + (2^{n-1} - 2) - (n-2) \\
&= 2^{n-1}
\end{aligned}$$

Chain Queries, Zig-Zag Join Trees, No Cartesian Product All possible zig-zag trees can be derived from a left-deep tree by exchanging the left and right arguments of a subset of the joins. Since for the first join these alternatives are already considered within the set of left-deep trees, we are left with $n - 2$ joins. Hence, the number of zig-zag trees for n relations in a chain query is $2^{n-2} * 2^{n-1} = 2^{2n-3}$.

Chain Queries, Bushy Join Trees, No Cartesian Product We can compute the number of bushy trees with no cross products for a chain query in the following way. Let us denote this number by $f(n)$. Again, let us assume that the chain query has the form $R_1 - R_2 - \dots - R_{n-1} - R_n$. For $n = 0$, we only have the empty join tree. For $n = 1$ we have one join tree. For $n = 2$ we have two join trees. For more relations, every subtree of the join tree must contain a subchain in order to avoid cross products. Further, the subchain can occur

as the left or right argument of the join. Hence, we can compute $f(n)$ as

$$\sum_{k=1}^{n-1} 2 f(k) f(n-k)$$

This is equal to

$$2^{n-1} \mathcal{C}(n-1)$$

EX

where $\mathcal{C}(n)$ are the Catalan Numbers.

Star Queries, No Cartesian Product The first join has to connect the center relation R_0 with any of the other relations. The other relations can follow in any order. Since R_0 can be the left or the right input of the first join, there are $2 * (n-1)!$ possible left-deep join trees for Star Queries with no Cartesian Product.

The number of zig-zag join trees is derived by exchanging the arguments of all but the first join in any left-deep join tree. We cannot consider the first join since we did so in counting left-deep join trees. Hence, the total number of zig-zag join trees is $2 * (n-1)! * 2^{n-2} = 2^{n-1} * (n-1)!$.

Constructing bushy join trees with no Cartesian Product from a Star Query other than zig-zag join trees is not possible.

Remarks The numbers for star queries are also upper bounds for tree queries. For clique queries, no join tree containing a cross product is possible. Hence, all join trees are valid join trees and the search space size is the same as the corresponding search space for join trees with cross products.

To give the reader a feeling for the numbers, the following tables contain the potential search space sizes for some n .

Join trees without cross products					
		chain query		star query	
	left-deep	zig-zag	bushy	left-deep	zig-zag/bushy
n	2^{n-1}	2^{2n-3}	$2^{n-1}\mathcal{C}(n-1)$	$2 * (n-1)!$	$2^{n-1}(n-1)!$
1	1	1	1	1	1
2	2	2	2	2	2
3	4	8	8	4	8
4	8	32	40	12	48
5	16	128	224	48	384
6	32	512	1344	240	3840
7	64	2048	8448	1440	46080
8	128	8192	54912	10080	645120
9	256	32768	366080	80640	10321920
10	512	131072	2489344	725760	185794560

	With cross products/cliقة		
	left-deep	zig-zag	bushy
n	$n!$	$2^{n-2} * n!$	$n!C(n-1)$
1	1	1	1
2	2	2	2
3	6	12	12
4	24	96	120
5	120	960	1680
6	720	11520	30240
7	5040	161280	665280
8	40320	2580480	17297280
9	362880	46448640	518918400
10	3628800	928972800	17643225600

Note that in Figure 2.6 only 32 join trees are listed, whereas the number of bushy trees for chain queries with four relations is 40. The missing eight cases are those zig-zag trees which are symmetric (i.e. derived by applying commutativity to all occurring joins) to the ones contained in the second column.

From these numbers, it becomes immediately clear why historically the search space of query optimizers was restricted to left-deep trees and cross products for connected query graphs were not considered.

3.1.6 Problem Complexity

The complexity of the join ordering problem depends on several parameters. These are the shape of the query graph, the class of join trees to be considered, whether cross products are considered or not, and whether the cost function has the ASI property or not. Not for all the combinations complexity results are known. What is known is summarized in the following table.

Query graph	Join tree	Cross products	Cost function	Complexity
general	left-deep	no	ASI	NP-hard
tree/star/chain	left-deep	no	one join method (ASI)	P
star	left-deep	no	two join methods (NLJ+SMJ)	NP-hard
general/tree/star	left-deep	yes	ASI	NP-hard
chain	left-deep	yes	—	open
general	bushy	no	ASI	NP-hard
tree	bushy	no	—	open
star	bushy	no	ASI	P
chain	bushy	no	any	P
general	bushy	yes	ASI	NP-hard
tree/star/chain	bushy	yes	ASI	NP-hard

Ibaraki and Kameda were the first who showed that the problem of deriving optimal left-deep trees for cyclic queries is NP-hard for a cost function for an n-way nested loop join implementation [377]. The proof was repeated for the cost function C_{out} which has the ASI property [165, 760]. In both proofs, the

clique problem was used for the reduction [273]. C_{out} was also used in the other proofs of NP-hardness results. The next line goes back to the same paper. Ibaraki and Kameda also described an algorithm to solve the join ordering problem for tree queries producing optimal left-deep trees for a special cost function for a nested-loop n-way join algorithm. Their algorithm was based on the observation that their cost function has the ASI property. For this case, they could derive an algorithm from an algorithm for a sequencing problem for job scheduling designed by Monma and Sidney [533]. They, in turn, used an earlier result by Lawler [459]. The algorithm of Ibaraki and Kameda was slightly generalized by Krishnamurthy, Boral, and Zaniolo, who were also able to sketch a more efficient algorithm. It improves the time bounds from $O(n^2 \log n)$ to $O(n^2)$. The disadvantage of both approaches is that with every relation, a fixed (i.e. join-tree independent) join implementation must be associated before the optimization starts. Hence, it only produces optimal trees if there is only one join implementation available or one is able to guess the optimal join method before hand. This might not be the case. The polynomial algorithm which we term IKKBZ is described in Section 3.2.2.

For star queries, Ganguly investigated the problem of generating optimal left-deep trees if no cross products but two different cost functions (one for nested loop join, the other for sort merge join) are allowed. It turned out that this problem is NP-hard [267].

The next line is due to Cluet and Moerkotte [165]. They showed by reduction from 3DM that taking into account cross products results in an NP-hard problem even for star queries. Remember that star queries are tree queries and general graphs.

The problem for general bushy trees follows from a result by Scheufele and Moerkotte [657]. They showed that building optimal bushy trees for cross products only (i.e. all selectivities equal one) is already NP-hard. This result also explains the last two lines.

By noting that for star queries, all bushy trees that do not contain a cross product are left-deep trees, the problem can be solved by the IKKBZ algorithm for left-deep trees. Ono and Lohman showed that for chain queries dynamic programming considers only a polynomial number of bushy trees if no cross products are considered [553]. This is discussed in Section 3.2.4.

The table is rather incomplete. Many open problems exist. For example, if we have chain queries and consider cross products: is the problem NP-hard or in P? Some results for this problem have been presented [657], but it is still an open problem (see Section 3.2.7). Open is also the case where we produce optimal bushy trees with no cross products for tree queries. Yet another example of an open problem is whether we could drop the ASI property and are still able to derive a polynomial algorithm for a tree query. This is especially important, since the cost function for a sort-merge algorithm does not have the ASI property.

Good summaries of complexity results for different join ordering problems can be found in the theses of Scheufele [655] and Hamalainen [338].

Given that join ordering is an inherently complex problem with no polynomial algorithm in sight, one might wonder whether there exists good polynomial

approximation algorithms. Chances are that even this is not the case. Chatterji, Evani, Ganguly, and Yemmanuru showed that three different optimization problems — all asking for linear join trees — are not approximable [119].

3.2 Deterministic Algorithms

3.2.1 Heuristics

We now present some simple heuristic solutions to the problem of join ordering. These heuristics only produce left-deep trees. Since left-deep trees are equivalent with permutations, these heuristics order the joins according to some criterion.

The core algorithm for the heuristics discussed here is the *greedy algorithm* (for an introduction see [179]). In greedy algorithms, a *weight* is associated with each entity. In our case, weights are associated with each relation. A typical weight function is the cardinality of the relation ($|R|$). Given a weight function *weight*, a greedy join ordering algorithm works as follows:

```
GreedyJoinOrdering-1( $\{R_1, \dots, R_n\}$ , (*weight)(Relation))
Input: a set of relations to be joined and a weight function
Output: a join order
 $S = \epsilon$ ; // initialize  $S$  to the empty sequence
 $R = \{R_1, \dots, R_n\}$ ; // let  $R$  be the set of all relations
while(!empty( $R$ )) {
  Let  $k$  be such that:  $\text{weight}(R_k) = \min_{R_i \in R}(\text{weight}(R_i))$ ;
   $R \setminus = R_k$ ; // eliminate  $R_k$  from  $R$ 
   $S \circ = R_k$ ; // append  $R_k$  to  $S$ 
}
return  $S$ 
```

This algorithm takes cross products into account. If we are only interested in left-deep join trees with no cross products, we have to require that R_k is connected to some of the relations contained in S in case $S \neq \epsilon$. Note that a more efficient implementation would sort the relations according to their weight.

Not all heuristics can be implemented with a greedy algorithm as simple as above. An often-used heuristics is to take the relation next that produces the smallest (next) intermediate result. This cannot be determined by the relation alone. One must take into account the sequence S already processed, since only then the selectivities of all predicates connecting relations in S and the new relation are deducible. And we must take the product of these selectivities and the cardinality of the new relation in order to get an estimate of the intermediate result's cardinality. As a consequence, the weights become *relative* to S . In other words, the weight function now has two parameters: a sequence of relations already joined and the relation whose relative weight is to be computed. Here is the next algorithm:


```

GreedyJoinOrdering-2( $\{R_1, \dots, R_n\}$ ,
    (*weight)(Sequence of Relations, Relation))
Input: a set of relations to be joined and a weight function
Output: a join order
 $S = \epsilon$ ; // initialize  $S$  to the empty sequence
 $R = \{R_1, \dots, R_n\}$ ; // let  $R$  be the set of all relations
while(!empty( $R$ )) {
    Let  $k$  be such that:  $\text{weight}(S, R_k) = \min_{R_i \in R}(\text{weight}(S, R_i))$ ;
     $R \setminus = R_k$ ; // eliminate  $R_k$  from  $R$ 
     $S \circ = R_k$ ; // append  $R_k$  to  $S$ 
}
return  $S$ 

```

Note that for this algorithm, sorting is not possible. GreedyJoinOrdering-2 can be improved by taking every relation as the starting one.

```

GreedyJoinOrdering-3( $\{R_1, \dots, R_n\}$ , (*weight)(Sequence of Relations, Relation))
Input: a set of relations to be joined and a weight function
Output: a join order
Solutions =  $\emptyset$ ;
for ( $i = 1$ ;  $i \leq n$ ;  $++i$ ) {
     $S = R_i$ ; // initialize  $S$  to a singleton sequence
     $R = \{R_1, \dots, R_n\} \setminus \{R_i\}$ ; // let  $R$  be the set of all relations
    while(!empty( $R$ )) {
        Let  $k$  be such that:  $\text{weight}(S, R_k) = \min_{R_i \in R}(\text{weight}(S, R_i))$ ;
         $R \setminus = R_k$ ; // eliminate  $R_k$  from  $R$ 
         $S \circ = R_k$ ; // append  $R_k$  to  $S$ 
    }
    Solutions +=  $S$ ;
}
return cheapest in Solutions

```

In addition to the relative weight function mentioned before, another often used relative weight function is the product of the selectivities connecting relations in S with the new relation. This heuristic is sometimes called *MinSel*.

The above two algorithms generate linear join trees. Fegaras proposed a heuristic (named Greedy Operator Ordering (GOO)) to generate bushy join trees [236, 237]. The idea is as follows. A set of join trees **Trees** is initialized such that it contains all the relations to be joined. It then investigates all pairs of trees contained in **Trees**. Among all of these, the algorithm joins the two trees that result in the smallest intermediate result when joined. The two trees are then eliminated from **Trees** and the new join tree joining them is added to it. The algorithm then looks as follows:

```

GOO( $\{R_1, \dots, R_n\}$ )
Input: a set of relations to be joined
Output: join tree
Trees :=  $\{R_1, \dots, R_n\}$ 
while ( $|\text{Trees}| \neq 1$ ) {
    find  $T_i, T_j \in \text{Trees}$  such that  $i \neq j$ ,  $|T_i \bowtie T_j|$  is minimal
        among all pairs of trees in Trees
    Trees - =  $T_i$ ;
    Trees - =  $T_j$ ;
    Trees + =  $T_i \bowtie T_j$ ;
}
return the tree contained in Trees;

```

Our GOO variant differs slightly from the one proposed by Fegaras. He uses arrays, explicitly handles the forming of the join predicates, and materializes intermediate result sizes. Hence, his algorithm is a little more elaborated, but we assume that the reader is able to fill in the gaps.

None of our algorithms so far considers different join implementations. An explicit consideration of commutativity for non-symmetric cost functions could also help to produce better join trees. The reader is asked to work out the details of these extensions. In general, the heuristics do not produce the optimal plan. EX
The reader is advised to find examples where the heuristics are far off the best possible plan. EX

3.2.2 Determining the Optimal Join Order in Polynomial Time

Since the general problem of join ordering is NP-hard, we cannot expect to find a polynomial solution for it. However, for special cases, we can expect to find solutions that work in polynomial time. These solutions can also be used as heuristics for the general case, either to find a not-that-bad join tree or to determine an upper bound for the costs that is then fed into a search procedure to prune the search space.

The most general case for which a polynomial solution is known is characterized by the following features:

- the query graph must be acyclic
- no cross products are considered
- the search space is restricted to left-deep trees
- the cost function must have the ASI property

The algorithm was presented by Ibaraki and Kameda [377]. Later Krishnamurthy, Boral, and Zaniolo presented it again for some other cost functions (still having the ASI property) [443]. They also observed that the upper bound $O(n^2 \log n)$ of the original algorithm could be improved to $O(n^2)$. In any case, the algorithm is based on an algorithm discovered by Monma and Sidney for job scheduling [459, 533]. Let us call the (unimproved) algorithm IKKBZ-Algorithm.

The IKKBZ-Algorithm considers only join operations that have a cost function of the form:

$$\text{cost}(R_i \bowtie R_j) = |R_i| * h_j(|R_j|)$$

where each R_j can have its own cost function h_j . We denote the set of h_j by H and parameterize cost functions with it. Example instantiations are

- $h_j \equiv 1.2$ for main memory hash-based joins
- $h_j \equiv \text{id}$ for nested-loop joins

where id is the identity function. Let us denote by n_i the cardinality of the relation R_i ($n_i := |R_i|$). Then, the $h_i(n_i)$ represent the costs per input tuple to be joined with R_i .

The algorithm works as follows. For every relation R_k it computes the optimal join order under the assumption that R_k is the first relation in the join sequence. The resulting subproblems then resemble a job-scheduling problem that can be solved by the Monma-Sidney-Algorithm [533].

In order to present this algorithm, we need the notion of a *precedence graph*. A *precedence graph* is formed by taking a node in the (undirected) query graph and making this node a root node of a (directed) tree where the edges point away from the selected root node. Hence, for acyclic, connected query graphs—those we consider in this section—a precedence graph is a tree. We construct the precedence graph of a query graph $G = (V, E)$ as follows:


- Make some relation $R_k \in V$ the root node of the precedence graph.
- As long as not all relations are included in the precedence graph: Choose a relation $R_i \in V$, such that $(R_j, R_i) \in E$ is an edge in the query graph and R_j is already contained in the (partial) precedence graph constructed so far and R_i is not. Add R_j and the edge $R_j \rightarrow R_i$ to the precedence graph.

A sequence $S = v_1, \dots, v_k$ of nodes conforms to a precedence graph $G = (V, E)$ if the following conditions are satisfied:


1. for all i ($2 \leq i \leq k$) there exists a j ($1 \leq j < i$) with $(v_j, v_i) \in E$ and
2. there is no edge $(v_i, v_j) \in E$ for $i > j$.

For non-empty sequences U and V in a precedence graph, we write $U \rightarrow V$ if, according to the precedence graph, U must occur before V . This requires U and V to be disjoint. More precisely, there can only be paths from nodes in U to nodes in V and at least one such path exists.


Consider the following query graph:



For this query graph, we can derive the following precedence graphs:



The IKKBZ-Algorithm takes a single precedence graph and produces a new one that is totally ordered. From this total order it is very easy to construct a corresponding join graph: the following figure contains a totally ordered precedence graph (left-hand side) as generated by the IKKBZ-Algorithm and the corresponding join graph on the right-hand side.



Define

$$R_{1,2,\dots,k} := R_1 \bowtie R_2 \bowtie \dots \bowtie R_k$$

$$n_{1,2,\dots,k} := |R_{1,2,\dots,k}|$$

For a given precedence graph, let R_i be a relation and \mathcal{R}_i be the set of relations from which there exists a path to R_i . Then, in any join tree adhering to the precedence graph, all relations in \mathcal{R}_i and only those will be joined before R_i . Hence, we can define $s_i = \prod_{R_j \in \mathcal{R}_i} f_{i,j}$ for $i > 1$. Note that for any i only one j

with $f_{i,j} \neq 1$ exists in the product. If the precedence graph is a chain, then the following holds:

$$n_{1,2,\dots,k+1} = n_{1,2,\dots,k} * s_{k+1} * n_{k+1}$$

We define $s_1 = 1$. Then we have

$$n_{1,2} = s_2 * (n_1 * n_2) = (s_1 * s_2) * (n_1 * n_2)$$

and, in general,

$$n_{1,2,\dots,k} = \prod_{i=1}^k (s_i * n_i).$$

We call the s_i selectivities, although they depend on the precedence graph.

The costs for a totally ordered precedence graph G can thus be computed as follows:

$$\begin{aligned} Cost_H(G) &= \sum_{i=2}^n [n_{1,2,\dots,i-1} * h_i(n_i)] \\ &= \sum_{i=2}^n [(\prod_{j=1}^{i-1} s_j * n_j) * h_i(n_i)] \end{aligned}$$

If we define $h_i(n_i) = s_i n_i$, then $Cost_H \equiv C_{out}$. The factor $s_i n_i$ determines by how much the input relation to be joined with R_i changes its cardinality after the join has been performed. If $s_i n_i$ is less than one, we call the join *decreasing*, if it is larger than one, we call the join *increasing*. This distinction plays an important role in the heuristic discussed in Section 3.2.3.

The cost function can also be defined recursively.

Definition 3.2.1 Define the cost function C_H as follows:

$$\begin{aligned} C_H(\epsilon) &= 0 \\ C_H(R_j) &= 0 \quad \text{if } R_j \text{ is the root} \\ C_H(R_j) &= h_j(n_j) \quad \text{else} \\ C_H(S_1 S_2) &= C_H(S_1) + T(S_1) * C_H(S_2) \end{aligned}$$

where

$$\begin{aligned} T(\epsilon) &= 1 \\ T(S) &= \prod_{R_i \in S} (s_i * n_i) \end{aligned}$$

It is easy to prove by induction that C_H is well-defined and that $C_H(G) = Cost_H(G)$.

Definition 3.2.2 Let A and B be two sequences and V and U two non-empty sequences. We say that a cost function C has the adjacent sequence interchange property (ASI property) if and only if there exists a function T and a rank function defined for sequences S as

$$\text{rank}(S) = \frac{T(S) - 1}{C(S)}$$

such that for non-empty sequences $S = AUVB$ the following holds

$$C(AUVB) \leq C(AVUB) \quad \Leftrightarrow \quad \text{rank}(U) \leq \text{rank}(V) \quad (3.1)$$

if $AUVB$ and $AVUB$ satisfy the precedence constraints imposed by a given precedence graph.

Lemma 3.2.3 The cost function C_H defined in Definition 3.2.1 has the ASI property.

The proof is very simple. Using the definition of C_H , we have

$$\begin{aligned} C_H(AUVB) &= C_H(A) \\ &\quad + T(A)C_H(U) \\ &\quad + T(A)T(U)C_H(V) \\ &\quad + T(A)T(U)T(V)C_H(B) \end{aligned}$$

and, hence,

$$\begin{aligned} C_H(AUVB) - C_H(AVUB) &= T(A)[C_H(V)(T(U) - 1) - C_H(U)(T(V) - 1)] \\ &= T(A)C_H(U)C_H(V)[\text{rank}(U) - \text{rank}(V)] \end{aligned}$$

The proposition follows. \square

Definition 3.2.4 Let $M = \{A_1, \dots, A_n\}$ be a set of node sequences in a given precedence graph. Then, M is called a module if for all sequences B that do not overlap with the sequences in M one of the following conditions holds:

- $B \rightarrow A_i, \forall 1 \leq i \leq n$
- $A_i \rightarrow B, \forall 1 \leq i \leq n$
- $B \not\rightarrow A_i$ and $A_i \not\rightarrow B, \forall 1 \leq i \leq n$

Lemma 3.2.5 Let C be any cost function with the ASI property and $\{A, B\}$ a module. If $A \rightarrow B$ and additionally $\text{rank}(B) \leq \text{rank}(A)$, then we find an optimal sequence among those in which B directly follows A .

Proof Every optimal permutation must have the form (U, A, V, B, W) , since $A \rightarrow B$. Assumption: $V \neq \epsilon$. If $\text{rank}(V) \leq \text{rank}(A)$, then we can exchange V and A without increasing the costs. If $\text{rank}(A) \leq \text{rank}(V)$, we have $\text{rank}(B) \leq \text{rank}(V)$ due to the transitivity of \leq . Hence, we can exchange B and V without increasing the costs. Both exchanges produce legal sequences obeying the precedence graph, since $\{A, B\}$ is a module. \square

If the precedence graph demands $A \rightarrow B$ but $\text{rank}(B) \leq \text{rank}(A)$, we speak of *contradictory sequences* A and B . Since the lemma shows that no non-empty subsequence can occur between A and B , we will combine A and B into a new single node replacing A and B . This node represents a *compound relation* comprising all relations in A and B . Its cardinality is computed by multiplying the cardinalities of all relations occurring in A and B , and its selectivity s is the product of all the selectivities s_i of the relations R_i contained in A and B . The continued process of this step until no more contradictory sequence exists is called *normalization*. The opposite step, replacing a compound node by the sequence of relations it was derived from, is called *denormalization*.

We can now present the algorithm IKKBZ.

IKKBZ(G)

Input: an acyclic query graph G for relations R_1, \dots, R_n

Output: the best left-deep tree

$R = \emptyset$;

for ($i = 1$; $i \leq n$; $++ i$) {

 Let G_i be the precedence graph derived from G and rooted at R_i ;

$T = \text{IKKBZ-Sub}(G_i)$;

$R+ = T$;

}

return best of R ;

IKKBZ-Sub(G_i)

Input: a precedence graph G_i for relations R_1, \dots, R_n rooted at some R_i

Output: the optimal left-deep tree under G_i

while (G_i is not a chain) {

 let r be the root of a subtree in G_i whose subtrees are chains;

 IKKBZ-Normalize(r);

 merge the chains under r according to the rank function

 in ascending order;

}

IKKBZ-Denormalize(G_i);

return G_i ;

IKKBZ-Normalize(r)

Input: the root r of a subtree T of a precedence graph $G = (V, E)$

Output: a normalized subchain

while ($\exists r', c \in V, r \rightarrow^* r', (r', c) \in E: \text{rank}(r') > \text{rank}(c)$) {

 replace r' by a compound relation r'' that represents $r'c$;

};

We do not give the details of IKKBZ-Denormalize, as it is trivial.




Figure 3.3: Illustrations for the IKKBZ Algorithm

Let us illustrate the algorithm IKKBZ-Sub by a simple example. We use the cost function C_{out} . Figure 3.3 A) shows a query graph. The relations are annotated with their sizes and the edges with the join selectivities. Choosing R_1 as the root of the precedence graph results in B). There, the nodes are annotated by the ranks of the relations. R_4 is the root of a subtree all of whose subtrees are chains. Hence, we normalize it. For R_5 , there is nothing to do. The ranks of R_6 and R_7 are contradictory. We form a compound relation $R_{6,7}$, calculate its cardinality, selectivity, and rank. The latter is shown in C). Merging the two subchains under R_4 results in D). Now R_1 is the root of a subtree with only chains underneath. Normalization detects that the ranks for R_4 and R_5 are contradictory. E) shows the tree after introducing the compound relation $R_{4,5}$. Now $R_{4,5}$ and $R_{6,7}$ have contradictory ranks, and we replace them by the compound relation $R_{4,5,6,7}$ as shown in F). Merging the chains under R_1 gives G). Since this is a chain, we leave the loop and denormalize. The final result is shown in H).




Figure 3.4: A query graph, its directed join graph, some spanning trees and join trees

We can use the IKKBZ-Algorithm to derive a heuristics also for cyclic queries, i.e. for general query graphs. In a first step, we determine a minimal spanning tree of the query graph. It is then used as the input query graph for the IKKBZ-Algorithm. Let us call this the *IKKBZ-based Heuristics*.

3.2.3 The Maximum-Value-Precedence Algorithm

Lee, Shih, and Chen proposed a very interesting heuristics for the join ordering problem [460]. They use a *weighted directed join graph* (WDJG) to represent queries. Within this graph, every join tree corresponds to a spanning tree. Given a conjunctive query with join predicates P . For a join predicate $p \in P$,

we denote by $\mathcal{R}(p)$ the relations whose attributes are mentioned in p .

Definition 3.2.6 *The directed join graph of a conjunctive query with join predicates P is a triple $G = (V, E_p, E_v)$, where V is the set of nodes and E_p and E_v are sets of directed edges defined as follows. For any two nodes $u, v \in V$, if $\mathcal{R}(u) \cap \mathcal{R}(v) \neq \emptyset$ then $(u, v) \in E_p$ and $(v, u) \in E_p$. If $\mathcal{R}(u) \cap \mathcal{R}(v) = \emptyset$, then $(u, v) \in E_v$ and $(v, u) \in E_v$. The edges in E_p are called physical edges, those in E_v virtual edges.*

Note that in G for every two nodes u, v , there is an edge (u, v) that is either physical or virtual. Hence, G is a clique.

Let us see how we can derive a join tree from a spanning tree of a directed join graph. Figure 3.4 I) gives a simple query graph Q corresponding to a chain and Part II) presents Q 's directed join graph. Physical edges are drawn by solid arrows, virtual edges by dotted arrows. Let us first consider the spanning tree shown in Part III a). It says that we first execute $R_1 \bowtie_{p_{1,2}} R_2$. The next join predicate to evaluate is $p_{2,3}$. Obviously, it does not make much sense to execute $R_2 \bowtie_{p_{2,3}} R_3$, since R_1 and R_2 have already been joined. Hence, we replace R_2 in the second join by the result of the first join. This results in the join tree $(R_1 \bowtie_{p_{1,2}} R_2) \bowtie_{p_{2,3}} R_3$. For the same reason, we proceed by joining this result with R_4 . The final join tree is shown in Part III b). Part IV a) shows another spanning tree. The two joins $R_1 \bowtie_{p_{1,2}} R_2$ and $R_3 \bowtie_{p_{3,4}} R_4$ can be executed independently and do not influence each other. Next, we have to consider $p_{2,3}$. Both R_2 and R_3 have already been joined. Hence, the last join processes both intermediate results. The final join tree is shown in Part IV b). The spanning tree shown in Part V a) results in the same join tree shown in Part V b). Hence, two different spanning trees can result in the same join tree. However, the spanning tree in Part IV a) is more specific in that it demands $R_1 \bowtie_{p_{1,2}} R_2$ to be executed before $R_3 \bowtie_{p_{3,4}}$.

Next, take a look at Figure 3.5. Part I), II), and III a) show a query graph, its directed join tree and a spanning tree. To build a join tree from the spanning tree we proceed as follows. We have to execute $R_2 \bowtie_{p_{2,3}} R_3$ and $R_3 \bowtie R_4$ first. In which way we do so is not really fixed by the spanning tree. So let us do both in parallel. Next is $p_{1,2}$. The only dependency the spanning tree gives us is that it should be executed after $p_{3,4}$. Since there is no common relation between those two, we perform $R_1 \bowtie_{p_{1,2}} R_2$. Last is $p_{4,5}$. Since we find $p_{3,4}$ below it, we use the intermediate result produced by it as a replacement for R_4 . The result is shown in Part III b). It has three loose ends. Additional joins are required to tie the partial results together. Obviously, this is not what we want. A spanning tree that avoids this problem of additional joins is called *effective*. It can be shown that a spanning tree $T = (V, E)$ is effective if it satisfies the following conditions [460]:

1. T is a binary tree,
2. for all inner nodes v and node u with $(u, v) \in E$ it holds that $\mathcal{R}^*(T(u)) \cap \mathcal{R}(v) \neq \emptyset$, and
3. for all nodes v, u_1, u_2 with $u_1 \neq u_2$, $(u_1, v) \in E$, and $(u_2, v) \in E$ one of the following two conditions holds:




Figure 3.5: A query graph, its directed join tree, a spanning tree and its problem

- (a) $((\mathcal{R}^*(T(u_1)) \cap \mathcal{R}(v)) \cap (\mathcal{R}^*(T(u_2)) \cap \mathcal{R}(v))) = \emptyset$ or
 (b) $(\mathcal{R}^*(T(u_1)) \cap \mathcal{R}(v) = \mathcal{R}(v)) \vee (\mathcal{R}^*(T(u_2)) \cap \mathcal{R}(v) = \mathcal{R}(v))$.

Thereby, we denote by $T(v)$ the partial tree rooted at v and by $\mathcal{R}^*(T') = \cup_{v \in T'} \mathcal{R}(v)$ the set of all relations in subtree T' .

We see that the spanning tree in Figure 3.5 III a) is ineffective since, for example, $\mathcal{R}(p_{2,3}) \cap \mathcal{R}(p_{4,5}) = \emptyset$. The spanning tree in Figure 3.4 IV a) is also ineffective. During the algorithm we will take care—by checking the above conditions—that only effective spanning trees are generated.

We now assign weights to the edges of the directed join graph. For two nodes $v, u \in V$ define $u \sqcap v := \mathcal{R}(u) \cap \mathcal{R}(v)$. For simplicity, we assume that every predicate involves exactly two relations. Then for all $u, v \in V$, $u \sqcap v$ contains a single relation. Let $v \in V$ be a node with $\mathcal{R}(v) = \{R_i, R_j\}$. We abbreviate $R_i \bowtie_v R_j$ by \bowtie_v . Using these notations, we can attach weights to the edges to define the *weighted directed join graph*.

Definition 3.2.7 Let $G = (V, E_p, E_v)$ be a directed join graph for a conjunctive query with join predicates P . The weighted directed join graph is derived from G by attaching a weight to each edge as follows:

- Let $(u, v) \in E_p$ be a physical edge. The weight $w_{u,v}$ of (u, v) is defined as

$$w_{u,v} = \frac{|\bowtie_u|}{|u \sqcap v|}.$$

- For virtual edges $(u, v) \in E_v$, we define $w_{u,v} = 1$.

(Lee, Shih, and Chen actually attach two weights to each edge: one additional weight for the size of the tuples (in bytes) [460].)

The weights of physical edges are equal to the s_i of the dependency graph used in the IKKBZ-Algorithm (Section 3.2.2). To see this, assume $\mathcal{R}(u) = \{R_1, R_2\}$, $\mathcal{R}(v) = \{R_2, R_3\}$. Then

$$\begin{aligned} w_{u,v} &= \frac{|\bowtie_u|}{|u \sqcap v|} \\ &= \frac{|R_1 \bowtie_u R_2|}{|R_2|} \\ &= \frac{f_{1,2} |R_1| |R_2|}{|R_2|} \\ &= f_{1,2} |R_1| \end{aligned}$$

Hence, if the join $R_1 \bowtie_u R_2$ is executed before the join $R_2 \bowtie_v R_3$, the input size to the latter join changes by a factor $w_{u,v}$. This way, the influence of a join on another join is captured by the weights. Since those nodes connected by a virtual edge do not influence each other, a weight of 1 is appropriate.

Additionally, we assign weights to the nodes of the directed join graph. The weight of a node reflects the change in cardinality to be expected when certain other joins have been executed before. They are specified by a (partial) spanning tree S . Given S , we denote by $\bowtie_{p_{i,j}}^S$ the result of the join $\bowtie_{p_{i,j}}$ if all joins preceding $p_{i,j}$ in S have been executed. Then the weight attached to node $p_{i,j}$ is defined as

$$w(p_{i,j}, S) = \frac{|\bowtie_{p_{i,j}}^S|}{|R_i \bowtie_{p_{i,j}} R_j|}.$$

For empty sequences ϵ , we define $w(p_{i,j}, \epsilon) = |R_i \bowtie_{p_{i,j}} R_j|$. Similarly, we define the cost of a node $p_{i,j}$ depending on other joins preceding it in some given spanning tree S . We denote this by $\text{cost}(p_{i,j}, S)$. The actual cost function can be one we have introduced so far or any other one. In fact, if we have a choice of several join implementations, we can take the minimum over all their cost functions. This then chooses the most effective join implementation.

The maximum value precedence algorithm works in two phases. In the first phase, it searches for edges with a weight smaller than one. Among these, the one with the biggest impact is chosen. This one is then added to the spanning tree. In other words, in this phase, the costs of expensive joins are minimized by making sure that (*size*) *decreasing joins* are executed first. The second phase adds edges such that the intermediate result sizes increase as little as possible.

MVP(G)

Input: a weighted directed join graph $G = (V, E_p, E_v)$

Output: an effective spanning tree

Q_1 .insert(V); /* priority queue with smallest node weights $w(\cdot)$ first */

$Q_2 = \emptyset$; /* priority queue with largest node weights $w(\cdot)$ first */

$G' = (V', E')$ with $V' = V$ and $E' = E_p$; /* working graph */

$S = (V_S, E_S)$ with $V_S = V$ and $E_S = \emptyset$; /* resulting effective spanning tree */

while (! Q_1 .empty()) && $|E_S| < |V| - 1$ { /* Phase I */

```

v = Q1.head();
among all (u, v) ∈ E', wu,v < 1 such that
  S' = (V, E'S) with E'S = ES ∪ {(u, v)} is acyclic and effective
  select one that maximizes cost(⋈v, S) - cost(⋈v, S');
if (no such edge exists) {
  Q1.remove(v);
  Q2.insert(v);
  continue;
}
MvpUpdate((u, v));
recompute w(·) for v and its ancestors; /* rearranges Q1 */
}
while (!Q2.empty() && |ES| < |V| - 1) { /* Phase II */
  v = Q2.head();
  among all (u, v), (v, u) ∈ E' denoted by (x, y) henceforth
  such that
    S' = (V, E'S) with E'S = ES ∪ {(x, y)} is acyclic and effective
    select the one that minimizes cost(⋈v, S') - cost(⋈v, S);
  MvpUpdate((x, y));
  recompute w(·) for y and its ancestors; /* rearranges Q2 */
}
return S;

MvpUpdate((u, v))
Input: an edge to be added to S
Output: side-effects on S, G',
  ES ∪ = {(u, v)};
  E' \ = {(u, v), (v, u)};
  E' \ = {(u, w) | (u, w) ∈ E'}; /* (1) */
  E' ∪ = {(v, w) | (u, w) ∈ Ep, (v, w) ∈ Ev}; /* (3) */
  if (v has two inflowing edges in S) { /* (2) */
    E' \ = {(w, v) | (w, v) ∈ E'};
  }
  if (v has one outflowing edge in S) { /* (1) in paper but not needed */
    E' \ = {(v, w) | (v, w) ∈ E'};
  }
}

```

Note that in order to test for the effectiveness of a spanning tree in the algorithm, we just have to check the conditions for the node the selected edge leads to.

MvpUpdate first adds the selected edge to the spanning tree. It then eliminates edges that need not to be considered for building an effective spanning tree. Since (u, v) has been added, both (u, v) and (v, u) do not have to be considered any longer. Also, since effective spanning trees are binary trees, (1) every node must have only one parent node and (2) at most two child nodes. The edges leading to a violation are eliminated by **MvpUpdate** in the lines com-

mented with the corresponding numbers. For the line commented (3) we have the situation that $u \rightarrow v \dashrightarrow w$ and $u \rightarrow w$ in G . This means that u and w have common relations, but v and w do not. Hence, the result of performing v on the result of u will have a common relation with w . Thus, we add a (physical) edge $v \rightarrow w$.

3.2.4 Dynamic Programming

Algorithms

Consider the two join trees

$$(((R_1 \bowtie R_2) \bowtie R_3) \bowtie R_4) \bowtie R_5$$

and

$$(((R_3 \bowtie R_1) \bowtie R_2) \bowtie R_4) \bowtie R_5.$$

If we know that $((R_1 \bowtie R_2) \bowtie R_3)$ is cheaper than $((R_3 \bowtie R_1) \bowtie R_2)$, we know that the first join tree is cheaper than the second. Hence, we could avoid generating the second alternative and still won't miss the optimal join tree. The general principle behind this is the *optimality principle* (see [178]). For the join ordering problem, it can be stated as follows.¹

Let T be an optimal join tree for relations R_1, \dots, R_n . Then, every subtree S of T must be an optimal join tree for the relations it contains.

To see why this holds, assume that the optimal join tree T for relations R_1, \dots, R_n contains a subtree S which is not optimal. That is, there exists another join tree S' for the relations contained in S with strictly lower costs. Denote by T' the join tree derived by replacing S in T by S' . Since S' contains the same relations as S , T' is a join tree for the relations R_1, \dots, R_n . The costs of the join operators in T and T' that are not contained in S and S' are the same. Then, since the total cost of a join tree is the sum of the costs of the join operators and S' has lower costs than S , T' has lower costs than T . This contradicts the optimality of T .

The idea of dynamic programming applied to the generation of optimal join trees now is to generate optimal join trees for subsets of R_1, \dots, R_n in a bottom-up fashion. First, optimal join trees for subsets of size one, i.e. single relations, are generated. From these, optimal join trees of size two, three and so on until n are generated.

Let us first consider generating optimal left-deep trees. There, join trees for subsets of size k are generated from subsets of size $k - 1$ by adding a new join operator whose left argument is a join tree for $k - 1$ relations and whose right argument is a single relation. Exchanging left and right gives us the procedure for generating right-deep trees. If we want to generate zig-zag trees since our cost function is asymmetric, we have to consider both alternatives and take the cheapest one. We capture this in a procedure `CreateJoinTree` that takes

¹The optimality principle does not hold in the presence of properties.

two join trees as arguments and generates the above-mentioned alternatives. In case we want to consider different implementations for the join, we have to perform the above steps for all of them and return the cheapest alternative. Summarizing, the pseudo-code for `CreateJoinTree` looks as follows:

```

CreateJoinTree( $T_1, T_2$ )
Input: two (optimal) join trees  $T_1$  and  $T_2$ .
           for linear trees, we assume that  $T_2$  is a single relation
Output: an (optimal) join tree for joining  $T_1$  and  $T_2$ .
BestTree = NULL;
for all implementations impl do {
  if(!RightDeepOnly) {
    Tree =  $T_1 \bowtie^{impl} T_2$ 
    if (BestTree == NULL || cost(BestTree) > cost(Tree)) {
      BestTree = Tree;
    }
  }
  if(!LeftDeepOnly) {
    Tree =  $T_2 \bowtie^{impl} T_1$ 
    if (BestTree == NULL || cost(BestTree) > cost(Tree)) {
      BestTree = Tree;
    }
  }
}
return BestTree;

```

The boolean variables `RightDeepOnly` and `LeftDeepOnly` are used to restrict the search space to right-deep trees and left-deep trees. If both are false, zig-zag trees are generated. However, `CreateJoinTree` also generates bushy trees, if none of the input trees is a single relation.

In case of linear trees, T_2 will be the single relation in all of our algorithms. `CreateJoinTree` should not copy T_1 or T_2 . Instead, the newly generated join trees should share T_1 and T_2 by using pointers. Further, the join trees generated do not really need to be generated except for the final (best) join tree: the cost functions should be implemented such that they can be evaluated if they are given the left and right argument of the join.

Using `CreateJoinTree`, we are now ready to present our first dynamic programming algorithm in pseudo-code.

```

DP-Linear-1( $\{R_1, \dots, R_n\}$ )
Input: a set of relations to be joined
Output: an optimal left-deep (right-deep, zig-zag) join tree
for (i = 1; i <= n; ++i) {
  BestTree( $\{R_i\}$ ) =  $R_i$ ;
}

```




Figure 3.6: Search space with sharing under optimality principle

```

}
for (i = 1; i < n; ++i) {
  for all  $S \subseteq \{R_1, \dots, R_n\}$ ,  $|S| = i$  do {
    for all  $R_j \in \{R_1, \dots, R_n\}$ ,  $R_j \notin S$  do {
      if (NoCrossProducts && !connected( $\{R_j\}, S$ )) {
        continue;
      }
      CurrTree = CreateJoinTree(BestTree( $S$ ),  $R_j$ );
       $S' = S \cup \{R_j\}$ ;
      if (BestTree( $S'$ ) == NULL || cost(BestTree( $S'$ )) > cost(CurrTree)) {
        BestTree( $S'$ ) = CurrTree;
      }
    }
  }
}
return BestTree( $\{R_1, \dots, R_n\}$ );

```

NoCrossProducts is a boolean variable indicating whether cross products should be investigated. Of course, if the join graph is not connected, there must be a cross product, but for DP-Linear-1 and subsequent algorithms we assume that it is connected. The boolean function `connected` returns true, if there is

XC search
space size
difference
problem

a join predicate between one of the relations in its first argument and one of the relations in its second. The variable `BestTree` keeps track of the best join trees generated for every subset of the relations $\{R_1, \dots, R_n\}$. How this is done may depend on several parameters. The approaches are to use a hash table or an array of size $2^n(-1)$. Another issue is how to represent the sets of relations. Typically, bitvector representations are used. Then, testing for membership, computing a set's complement, adding elements and unioning is cheap. Yet another issue is the order in which join trees are generated. The procedure `DP-Linear-1` takes the approach to generate the join trees for subsets of size $1, 2, \dots, n$. To do so, it must be able to access the subsets of $\{R_1, \dots, R_n\}$ or their respective join trees by their size. One possibility is to chain all the join trees for subsets of a given size k ($1 \leq k \leq n$) and to use an array of size n to keep pointers to the start of the lists. In this case, to every join tree the set of relations it contains is attached, in order to be able to perform the test $R_i \notin S$. One way to do this is to embed a bitvector into each join tree node.

Figure 3.6 illustrates how the procedure `DP-Linear-1` works. In its first loop, it initializes the bottom row of join trees of size one. Then it computes the join trees joining exactly two relations. This is indicated by the next group of join trees. Since the figure leaves out commutativity, only one alternative join tree for every subset of size two is generated. This changes for subsets of size three. There, three alternative join trees are generated. Only the best join tree is retained. This is indicated by the ovals that encircle three join trees. Only this best join tree of size three is used to generate the final best join tree.

The short clarification after the algorithm already adumbrated that the order in which join trees are generated is not compulsory. The only necessary condition is the following.

Let S be a subset of $\{R_1, \dots, R_n\}$. Then, before a join tree for S can be generated, the join trees for all relevant subsets of S must already be available.

Note that this formulation is general enough to also capture the generation of bushy trees. It is, however, a little vague due to its reference to “relevance”. For the different join tree classes, this term can be given a precise semantics.

EX

Let us take a look at an alternative order to join tree generation. Assume that sets of relations are represented as bitvectors. A bitvector is nothing more than a base two integer. Successive increments of an integer/bitvector lead to different subsets. Further, the above condition is satisfied. We illustrate this by a small example. Assume that we have three relations R_1, R_2, R_3 . The i -th bit from the right in a three-bit integer indicates the presence of R_i for $1 \leq i \leq 3$.

000	{ }
001	{ R_1 }
010	{ R_2 }
011	{ R_1, R_2 }
100	{ R_3 }
101	{ R_1, R_3 }
110	{ R_2, R_3 }
111	{ R_1, R_2, R_3 }

This observation leads to another formulation of our dynamic programming algorithm. For this algorithm, it is very convenient to use an array of size 2^n to represent $\text{BestTree}(S)$ for subsets S of $\{R_1, \dots, R_n\}$.

```

DP-Linear-2( $\{R_1, \dots, R_n\}$ )
Input: a set of relations to be joined
Output: an optimal left-deep (right-deep, zig-zag) join tree
for (i = 1; i <= n; ++i) {
    BestTree(1 << i) =  $R_i$ ;
}
for (S = 1; S <  $2^n$ ; ++S) {
    if (BestTree(S) != NULL) continue;
    for all  $i \in S$  do {
         $S' = S \setminus \{i\}$ ;
        CurrTree = CreateJoinTree(BestTree( $S'$ ),  $R_i$ );
        if (BestTree(S) == NULL || cost(BestTree(S)) > cost(CurrTree)) {
            BestTree(S) = CurrTree;
        }
    }
}
return BestTree( $2^n - 1$ );

```

DP-Linear-2 differs from DP-Linear-1 not only in the order in which join trees are generated. Another difference is that it takes cross products into account.

From DP-Linear-2, it is easy to derive an algorithm that explores the space of bushy trees.

```

DP-Bushy( $\{R_1, \dots, R_n\}$ )
Input: a set of relations to be joined
Output: an optimal bushy join tree
for (i = 1; i <= n; ++i) {
    BestTree(1 << i) =  $R_i$ ;
}
for (S = 1; S <  $2^n$ ; ++S) {
    if (BestTree(S) != NULL) continue;
    for all  $S_1 \subset S$  do {
         $S_2 = S \setminus S_1$ ;
        CurrTree = CreateJoinTree(BestTree( $S_1$ ), BestTree( $S_2$ ));
        if (BestTree(S) == NULL || cost(BestTree(S)) > cost(CurrTree)) {
            BestTree(S) = CurrTree;
        }
    }
}
return BestTree( $2^n - 1$ );

```

This algorithm also takes cross products into account. The critical part is the generation of all subsets of S . Fortunately, Vance and Maier [778] provide a code fragment with which subset bitvector representations can be generated very efficiently. In C, this fragment looks as follows:

```

S1 = S & - S;
do {
  /* do something with subset S1 */
  S1 = S & (S1 - S);
} while (S1 != S);

```

S represents the input set. S_1 iterates through all subsets of S where S itself and the empty set are not considered. Analogously, all supersets can be generated as follows:

```

S1 = ~S & - ~S;
/* do something with first superset S1 */
while (S1) {
  S1 = ~S & (S1 - ~S)
  /* do something with superset S1
}

```

S represents the input set. S_1 iterates through all supersets of S including S itself.

Excursion Problem: exploiting orderings devastates the optimality principle.
Example: ...

XC ToDo

Excursion Pruning ...

XC ToDo

Number of Entries to be stored in the dynamic programming table

If dynamic programming uses a static hash table, determining its size in advance is necessary as the search space sizes differ vastly for different query graphs. In general, for every connected subgraph of the query graph one entry must exist. Chains require far fewer entries than cliques. It would be helpful to have a small routine solving the following problem: given a query graph, how many connected subgraph are there? Unfortunately, this problem is #P-hard as Sutner, Satyanarayana, and Suffel showed [743]. They build on results by Valiant [776] and Lichtenstein [478]. (For a definition of #P-hard see the book by Lewis and Papadimitriou [476] or the original paper by Valiant [775].)

However, for specific cases, these numbers can be given. If cross products are considered, the number of join trees stored in the dynamic programming table is

$$2^n - 1$$

which is one for each non-empty subset of relations.

If we do not consider cross products, the number of entries in the dynamic programming table corresponds to the number of connected subgraphs of the query graph. For connected query graphs, we denote this by $\#csg$. For chains, cycles, stars, and cliques with n nodes, we have

$$\#csg^{\text{chain}}(n) = \frac{n(n+1)}{2} \quad (3.2)$$

$$\#csg^{\text{cycle}}(n) = n^2 - n + 1 \quad (3.3)$$

$$\#csg^{\text{star}}(n) = 2^{n-1} + n - 1 \quad (3.4)$$

$$\#csg^{\text{clique}}(n) = 2^n - 1 \quad (3.5)$$

These equations can be derived from the following by summing over $k > 1$ where k gives the size of the connected subset:

$$\begin{aligned} \#csg^{\text{chain}}(n, k) &= (n - k + 1) \\ \#csg^{\text{cycle}}(n, k) &= \begin{cases} 1 & n = k \\ n & \text{else} \end{cases} \\ \#csg^{\text{star}}(n, k) &= \begin{cases} n & k = 1 \\ \binom{n-1}{k-1} & k > 1 \end{cases} \\ \#csg^{\text{clique}}(n, k) &= \binom{n}{k} \end{aligned}$$

Number of Join Trees Investigated

The number of join trees investigated by dynamic programming was extensively studied by Ono and Lohman [552, 553]. In order to estimate these numbers, we assume that `CreateJoinTree` produces a single join tree and hence counts as one although it may evaluate the costs for several join alternatives. We further do not count the initial join trees containing only a single relation.

Join Trees With Cartesian Product For the analysis of dynamic programming variants that do consider cross products, the notion of *join-pair* is helpful. Let S_1 and S_2 be subsets of the nodes (relations) of the query graph. We say (S_1, S_2) is a *join-pair*, if and only if

1. S_1 and S_2 are disjoint

If (S_1, S_2) is a join-pair, then (S_2, S_1) is a join pair. Further, if T_1 is a join tree for the relations in S_1 and T_2 is one for those in S_2 , then we can construct two valid join trees $T_1 \bowtie T_2$ and $T_2 \bowtie T_1$ where the joins may be cross products. Hence, the number of join-pairs coincides with the search space a dynamic programming algorithm explores. In fact, the number of join-pairs is the minimum number of join trees any dynamic programming algorithm that considers cross products has to investigate.

If `CreateJoinTree` considers commutativity of joins, the number of calls to it is precisely expressed by the count of non-symmetric join-pairs. In other

implementations `CreateJoinTree` might be called for all join-pairs and, thus, may not consider commutativity. The two formulas below only count non-symmetric join pairs.

The numbers of linear and bushy join trees with cartesian product is easiest to determine. They are independent of the query graph. For linear join trees, the number of join trees investigated by dynamic programming is equal to the number of non-symmetric join-pairs which is

$$n2^{n-1} - \frac{n(n+1)}{2}$$

Dynamic programming investigates the following number of bushy trees if cross products are considered.

$$\frac{(3^n - 2^{n+1} + 1)}{2}$$

This is equal to the number of non-symmetric join-pairs.

Join Trees without Cross Products In this paragraph, we assume that the query graph is connected. For the analysis of dynamic programming variants that do not consider cross products, it is helpful to have the notion of a *csg-cmp-pair*. Let S_1 and S_2 be subsets of the nodes (relations) of the query graph. We say (S_1, S_2) is a *csg-cmp-pair*, if and only if

1. S_1 induces a connected subgraph of the query graph,
2. S_2 induces a connected subgraph of the query graph,
3. S_1 and S_2 are disjoint, and
4. there exists at least one edge connected a node in S_1 to a node in S_2 .

If (S_1, S_2) is a *csg-cmp-pair*, then (S_2, S_1) is a valid *csg-cmp-pair*. Further, if T_1 is a join tree for the relations in S_1 and T_2 is one for those in S_2 , then we can construct two valid join trees $T_1 \bowtie T_2$ and $T_2 \bowtie T_1$. Hence, the number of *csg-cmp-pairs* coincides with the search space a dynamic programming algorithm explores. In fact, the number of *csg-cmp-pairs* is the minimum number of join trees any dynamic programming algorithm that does not consider cross products has to investigate.

If `CreateJoinTree` considers commutativity of joins, the number of calls to it is precisely expressed by the count of non-symmetric *csg-cmp-pairs*. In other implementations `CreateJoinTree` might be called for all *csg-cmp-pairs* and, thus, may not consider commutativity.

Let us denote the number of non-symmetric *csg-cmp-pairs* by $\#ccp$. Then

$$\begin{aligned} \#ccp^{\text{chain}}(n) &= \frac{1}{6}(n^3 - n) \\ \#ccp^{\text{cycle}}(n) &= (n^3 - 2n^2 + n)/2 \\ \#ccp^{\text{star}}(n) &= (n-1)2^{n-2} \\ \#ccp^{\text{clique}}(n) &= (3^n - 2^{n+1} + 1)/2 \end{aligned}$$

These numbers have to be multiplied by two if we want to count all csg-cmp-pairs.

If we do not consider composite inners, that is we restrict ourselves to left-deep join trees, then dynamic programming makes the following number of calls to `CreateJoinTree` for chain queries [553]:

$$(n - 1)^2$$

The following table presents some results for the above formulas.

	without cross products			with cross products	
	chain		star	any query graph	
	linear	bushy	linear	linear	bushy
n	$(n - 1)^2$	$(n^3 - n)/6$	$(n - 1)2^{n-2}$	$n2^{n-1} - n(n + 1)/2$	$(3^n - 2^{n+1} + 1)/2$
2	1	1	1	1	1
3	4	4	4	6	6
4	9	10	12	22	25
5	16	20	32	65	90
6	25	35	80	171	301
7	36	56	192	420	966
8	49	84	448	988	3025
9	64	120	1024	2259	9330
10	81	165	2304	5065	28501

Compare this table with the actual sizes of the search spaces in Section 3.1.5.

The dynamic programming algorithms can be implemented very efficiently and often form the core of commercial plan generators. However, they have the disadvantage that no plan is generated if they run out of time or space since the search space they have to explore is too big. One possible remedy goes as follows. Assume that a dynamic programming algorithm is stopped in the middle of its way through its actual search space. Further assume that the largest plans generated so far involve k relations. Then the cheapest of the plans with k relations is completed by applying any heuristics (e.g. `MinSel`). The completed plan is then returned. In Section 3.4.5, we will see two alternative solutions. Another solution is presented in [417].

Generating Bushy Trees without Cross Products

We now discuss dynamic programming algorithms to generate bushy trees without cross products. For this section, we assume that the query graph is connected. We will present three algorithms. The first algorithm (`DPsize`) generates its plans in increasing size of subplans and, hence, is a generalization of `DP-Linear-1`. The second algorithm (`DPsub`) generates its plans by considering plans subsets as does `DP-Linear-2`. An analysis of these two algorithms reveals that both are far away from the lower bound presented in the previous sections. Thus, a third algorithm (`DPccp`) which reaches this lower bound is presented. The results of this section are taken from [530, 529].

```

DPsize
Input: a connected query graph with relations  $R = \{R_0, \dots, R_{n-1}\}$ 
Output: an optimal bushy join tree without cross products
for all  $R_i \in R$  {
    BestPlan( $\{R_i\}$ ) =  $R_i$ ;
}
for all  $1 < s \leq n$  ascending // size of plan
for all  $1 \leq s_1 < s$  { // size of left subplan
     $s_2 = s - s_1$ ; // size of right subplan
    for all  $S_1 \subset R : |S_1| = s_1$ 
         $S_2 \subset R : |S_2| = s_2$  {
            ++InnerCounter;
            if ( $\emptyset \neq S_1 \cap S_2$ ) continue;
            if not ( $S_1$  connected to  $S_2$ ) continue;
            ++CsgCmpPairCounter;
             $p_1 = \text{BestPlan}(S_1)$ ;
             $p_2 = \text{BestPlan}(S_2)$ ;
            CurrPlan = CreateJoinTree( $p_1, p_2$ );
            if ( $\text{cost}(\text{BestPlan}(S_1 \cup S_2)) > \text{cost}(\text{CurrPlan})$ ) {
                BestPlan( $S_1 \cup S_2$ ) = CurrPlan;
            }
        }
    }
}
}
OnoLohmanCounter = CsgCmpPairCounter / 2;
return BestPlan( $\{R_0, \dots, R_{n-1}\}$ );

```

Figure 3.7: Algorithm DPsize

Size-based enumeration: DPsize In general, dynamic programming generates solutions for a larger problem in a bottom-up fashion by combining solutions for smaller problems. Taking this description literally, we can construct optimal plans of size n by joining plans P_1 and P_2 of size k and $n - k$. We just have to take care that (1) the sets of relations contained in P_1 and P_2 do not overlap, and (2) there is a join predicate connecting a relation P_1 with a relation in P_2 . After this remark, we are prepared to understand the pseudocode for algorithm DPsize (see Fig. 3.7). A table **BestPlan** associates with each set of relations the best plan found so far. The algorithm starts by initializing this table with plans of size one, i.e. single relations. After that, it constructs plans of increasing size (loop over s). Thereby, the first size considered is two, since plans of size one have already been constructed. Every plan joining n relations can be constructed by joining a plan containing s_1 relations with a plan containing s_2 relations. Thereby, $s_i > 0$ and $s_1 + s_2 = n$ must hold. Thus, the pseudocode loops over s_1 and sets s_2 accordingly. Since for every possible size there exist many plans, two more loops are necessary in order to loop over the plans of sizes s_1 and s_2 . Then, conditions (1) and (2) from above are tested. Only if their outcome is positive, we consider joining the plans p_1 and p_2 . The result is a plan **CurrPlan**. Let S be the relations contained in **CurrPlan**. If

BestPlan does not contain a plan for the relations in S or the one it contains is more expensive than **CurrPlan**, we register **CurrPlan** with **BestPlan**.

The algorithm **DPsize** can be made more efficient in case of $s_1 = s_2$. The algorithm as stated cycles through all plans p_1 joining s_1 relations. For each such plan, all plans p_2 of size s_2 are tested. Assume that plans of equal size are represented as a linked list. If $s_1 = s_2$, then it is possible to iterate through the list for retrieving all plans p_1 . For p_2 we consider the plans succeeding p_1 in the list. Thus, the complexity can be decreased from $s_1 * s_2$ to $s_1 * s_2 / 2$. The following formulas are valid only for the variant of **DPsize** where this optimization has been incorporated (see [529] for details).

If the counter **InnerCounter** is initialized with zero at the beginning of the algorithm **DPsize**, then we are able to derive analytically its value after **DPsize** terminates. Since this value of the inner counter depends on the query graph, we have to distinguish several cases. For chain, cycle, star, and clique queries, we denote by $I_{\text{DPsize}}^{\text{chain}}$, $I_{\text{DPsize}}^{\text{cycle}}$, $I_{\text{DPsize}}^{\text{star}}$, and $I_{\text{DPsize}}^{\text{clique}}$ the value of **InnerCounter** after termination of algorithm **DPsize**.

For chain queries, we then have: $I_{\text{DPsize}}^{\text{chain}}(n) =$

$$\begin{cases} 1/48(5n^4 + 6n^3 - 14n^2 - 12n) & n \text{ even} \\ 1/48(5n^4 + 6n^3 - 14n^2 - 6n + 11) & n \text{ odd} \end{cases}$$

For cycle queries, we have: $I_{\text{DPsize}}^{\text{cycle}}(n) =$

$$\begin{cases} \frac{1}{4}(n^4 - n^3 - n^2) & n \text{ even} \\ \frac{1}{4}(n^4 - n^3 - n^2 + n) & n \text{ odd} \end{cases}$$

For star queries, we have: $I_{\text{DPsize}}^{\text{star}}(n) =$

$$\begin{cases} 2^{2n-4} - 1/4 \binom{2(n-1)}{n-1} + q(n) & n \text{ even} \\ 2^{2n-4} - 1/4 \binom{2(n-1)}{n-1} + 1/4 \binom{n-1}{(n-1)/2} + q(n) & n \text{ odd} \end{cases}$$

with $q(n) = n2^{n-1} - 5 * 2^{n-3} + 1/2(n^2 - 5n + 4)$. For clique queries, we have: $I_{\text{DPsize}}^{\text{clique}}(n) =$

$$\begin{cases} 2^{2n-2} - 5 * 2^{n-2} + 1/4 \binom{2n}{n} - 1/4 \binom{n}{n/2} + 1 & n \text{ even} \\ 2^{2n-2} - 5 * 2^{n-2} + 1/4 \binom{2n}{n} + 1 & n \text{ odd} \end{cases}$$

Note that $\binom{2n}{n}$ is in the order of $\Theta(4^n / \sqrt{n})$.

Proofs of the above formulas as well as implementation details for the algorithm **DPsize** can be found in [529].

Subset-Driven Enumeration: DPsub Fig. 3.8 presents the pseudocode for the algorithm **DPsub**. The algorithm first initializes the table **BestPlan** with all possible plans containing a single relation. Then, the main loop starts. It iterates over all possible non-empty subsets of $\{R_0, \dots, R_{n-1}\}$ and constructs the best possible plan for each of them. The enumeration makes use of a bitvector representation of sets: The integer i induces the current subset S with


```

DPsub
Input: a connected query graph with relations  $R = \{R_0, \dots, R_{n-1}\}$ 
Output: an optimal bushy join tree
for all  $R_i \in R$  {
    BestPlan( $\{R_i\}$ ) =  $R_i$ ;
}
for  $1 \leq i < 2^n - 1$  ascending {
     $S = \{R_j \in R \mid (\lfloor i/2^j \rfloor \bmod 2) = 1\}$ 
    if not (connected  $S$ ) continue;    // *
    for all  $S_1 \subset S, S_1 \neq \emptyset$  do {
        ++InnerCounter;
         $S_2 = S \setminus S_1$ ;
        if ( $S_2 = \emptyset$ ) continue;
        if not (connected  $S_1$ ) continue;
        if not (connected  $S_2$ ) continue;
        if not ( $S_1$  connected to  $S_2$ ) continue;
        ++CsgCmpPairCounter;
         $p_1 = \text{BestPlan}(S_1)$ ;
         $p_2 = \text{BestPlan}(S_2)$ ;
        CurrPlan = CreateJoinTree( $p_1, p_2$ );
        if ( $\text{cost}(\text{BestPlan}(S)) > \text{cost}(\text{CurrPlan})$ ) {
            BestPlan( $S$ ) = CurrPlan;
        }
    }
}
OnoLohmanCounter = CsgCmpPairCounter / 2;
return BestPlan( $\{R_0, \dots, R_{n-1}\}$ );

```

Figure 3.8: Algorithm DPsub

its binary representation. Taken as bitvectors, the integers in the range from 1 to $2^n - 1$ exactly represent the set of all non-empty subsets of $\{R_0, \dots, R_{n-1}\}$, including the set itself. Further, by starting with 1 and incrementing by 1, the enumeration order is valid for dynamic programming: for every subset, all its subsets are generated before the subset itself.

This enumeration is very fast, since increment by one is a very fast operation. However, the relations contained in S may not induce a connected subgraph of the query graph. Therefore, we must test for connectedness. The goal of the next loop over all subsets of S is to find the best plan joining all the relations in S . Therefore, S_1 ranges over all non-empty, strict subsets of S . This can be done very efficiently by applying the code snippet of Vance and Maier [777, 778]. Then, the subset of relations contained in S but not in S_1 is assigned to S_2 . Clearly, S_1 and S_2 are disjoint. Hence, only connectedness tests have to be performed. Since we want to avoid cross products, S_1 and S_2 both must induce connected subgraphs of the query graph, and there must be a join predicate between a relation in S_1 and one in S_2 . If these conditions are fulfilled, we can construct a plan **CurrPlan** by joining the plans associated with S_1 and S_2 . If

BestPlan does not contain a plan for the relations in S or the one it contains is more expensive than **CurrPlan**, we register **CurrPlan** with **BestPlan**.

For chain, cycle, star, and clique queries, we denote by $I_{\text{DPsub}}^{\text{chain}}$, $I_{\text{DPsub}}^{\text{cycle}}$, $I_{\text{DPsub}}^{\text{star}}$, and $I_{\text{DPsub}}^{\text{clique}}$ the value of **InnerCounter** after termination of algorithm **DPsub**.

For chains, we have

$$I_{\text{DPsub}}^{\text{chain}}(n) = 2^{n+2} - n^n - 3n - 4 \quad (3.6)$$

For cycles, we have

$$I_{\text{DPsub}}^{\text{cycle}}(n) = n2^n + 2^n - 2n^2 - 2 \quad (3.7)$$

For stars, we have

$$I_{\text{DPsub}}^{\text{star}}(n) = 2 * 3^{n-1} - 2^n \quad (3.8)$$

For cliques, we have

$$I_{\text{DPsub}}^{\text{clique}}(n) = 3^n - 2^{n+1} + 1 \quad (3.9)$$

The number of failures for the additional check can easily be calculated as $2^n - \#\text{csg}(n) - 1$.

Sample numbers Fig. 3.9 contains tables with values produced by our formulas for input query graph sizes between 2 and 20. For different kinds of query graphs, it shows the number of **csg-cmp-pairs** ($\#\text{ccp}$), and the values for the inner counter after termination of **DPsize** and **DPsub** applied to the different query graphs.

Looking at these numbers, we observe the following:

- For chain and cycle queries, the **DPsize** soon becomes much faster than **DPsub**.
- For star and clique queries, the **DPsub** soon becomes much faster than **DPsize**.
- Except for clique queries, the number of **csg-cmp-pairs** is orders of magnitude less than the value of *InnerCounter* for all DP-variants.

From the latter observation we can conclude that in almost all cases the tests performed by both algorithms in their innermost loop fail. Both algorithms are far away from the theoretical lower bound given by $\#\text{ccp}$. This conclusion motivates us to derive a new algorithm whose *InnerCounter* value is equal to the number of **csg-cmp-pairs**.

Csg-cmp-pair enumeration-based algorithm: DPccp The algorithm **DPsub** solves the join ordering problem for a given subset S of relations by considering all pairs of disjoint subproblems which were already solved. Since the enumeration of subsets is very fast, this is a very efficient strategy if the search space is dense, e.g. for clique queries. However, if the search space is sparse, e.g. for

	Chain			Cycle		
n	#ccp	DPsub	DPsize	#ccp	DPsub	DPsize
2	1	2	1	1	2	1
5	20	84	73	40	140	120
10	165	3962	1135	405	11062	2225
15	560	130798	5628	1470	523836	11760
20	1330	4193840	17545	3610	22019294	37900
	Star			Clique		
n	#ccp	DPsub	DPsize	#ccp	DPsub	DPsize
2	1	2	1	1	2	1
5	32	130	110	90	180	280
10	2304	38342	57888	28501	57002	306991
15	114688	9533170	57305929	7141686	14283372	307173877
20	4980736	2323474358	59892991338	1742343625	3484687250	309338182241

Figure 3.9: Size of the search space for different graph structures

chain queries, the `DPsub` algorithm considers many subproblems which are not connected and, therefore, are not relevant for the solution, i.e. the tests in the innermost loop fail for the majority of cases. The main idea of our algorithm `DPccp` is that it only considers pairs of connected subproblems. More precisely, the algorithm considers exactly the `csg-cmp-pairs` of a graph.

Thus, our goal is to efficiently enumerate all `csg-cmp-pairs` (S_1, S_2) . Clearly, we want to enumerate every pair once and only once. Further, the enumeration must be performed in an order valid for dynamic programming. That is, whenever a pair (S_1, S_2) is generated, all non-empty subsets of S_1 and S_2 must have been generated before as a component of a pair. The last requirement is that the overhead for generating a single `csg-cmp-pair` must be constant or at most linear. This condition is necessary in order to beat `DPsize` and `DPsub`.

If we meet all these requirements, the algorithm `DPccp` is easily specified: iterate over all `csg-cmp-pairs` (S_1, S_2) and consider joining the best plans associated with them. Figure 3.10 shows the pseudocode. The first steps of an example enumeration are shown in Figure 3.11. Thick lines mark the connected subsets while thin lines mark possible join edges. Note that the algorithm explicitly exploits join commutativity. This is due to our enumeration algorithm developed below. If (S_1, S_2) is a `csg-cmp-pair`, then either (S_1, S_2) or (S_2, S_1) will be generated, but never both of them. An alternative is to modify `CreateJoinTree` to take care of commutativity.

We proceed as follows. Next we discuss an algorithm enumerating non-empty connected subsets S_1 of $\{R_0, \dots, R_{n-1}\}$. Then, we show how to enumerate the complements S_2 such that (S_1, S_2) is a `csg-cmp-pair`.

Let us start the exposition by fixing some notations. Let $G = (V, E)$ be an undirected graph. For a node $v \in V$ define the *neighborhood* $\mathcal{N}(v)$ of v as $\mathcal{N}(v) := \{v' \mid (v, v') \in E\}$. For a subset $S \subseteq V$ of V we define the *neighborhood* of S as $\mathcal{N}(S) := \cup_{v \in S} \mathcal{N}(v) \setminus S$. The neighborhood of a set of nodes thus consists

DPccp

Input: a connected query graph with relations $R = \{R_0, \dots, R_{n-1}\}$

Output: an optimal bushy join tree

```

for all  $R_i \in R$  {
    BestPlan( $\{R_i\}$ ) =  $R_i$ ;
}
for all csg-cmp-pairs  $(S_1, S_2), S = S_1 \cup S_2$  {
    ++InnerCounter;
    ++OnoLohmanCounter;
     $p_1 = \text{BestPlan}(S_1)$ ;
     $p_2 = \text{BestPlan}(S_2)$ ;
    CurrPlan = CreateJoinTree( $p_1, p_2$ );
    if ( $\text{cost}(\text{BestPlan}(S)) > \text{cost}(\text{CurrPlan})$ ) {
        BestPlan( $S$ ) = CurrPlan;
    }
    CurrPlan = CreateJoinTree( $p_2, p_1$ );
    if ( $\text{cost}(\text{BestPlan}(S)) > \text{cost}(\text{CurrPlan})$ ) {
        BestPlan( $S$ ) = CurrPlan;
    }
}
CsgCmpPairCounter = 2 * OnoLohmanCounter;
return BestPlan( $\{R_0, \dots, R_{n-1}\}$ );
    
```

Figure 3.10: Algorithm DPccp





Figure 3.11: Enumeration Example for DPccp

of all nodes reachable by a single edge. Note that for all $S, S' \subset V$ we have $\mathcal{N}(S \cup S') = (\mathcal{N}(S) \cup \mathcal{N}(S')) \setminus (S \cup S')$. This allows for an efficient bottom-up calculation of neighborhoods.

The following statement gives a hint on how to construct an enumeration procedure for connected subsets. Let S be a connected subset of an undirected graph G and S' be any subset of $\mathcal{N}(S)$. Then $S \cup S'$ is connected. As a consequence, a connected subset can be enlarged by adding any subset of its neighborhood.

We could generate all connected subsets as follows. For every node $v_i \in V$ we perform the following enumeration steps: First, we emit $\{v_i\}$ as a connected subset. Then, we expand $\{v_i\}$ by calling a routine that extends a given connected set to bigger connected sets. Let the routine be called with some connected set S . It then calculates the neighborhood $\mathcal{N}(S)$. For every non-empty subset

Figure 3.12: Sample graph to illustrate `EnumerateCsgRec`

$N \subseteq \mathcal{N}(S)$, it emits $S' = S \cup N$ as a further connected subset and recursively calls itself with S' . The problem with this routine is that it produces duplicates.

This is the point where the breadth-first numbering comes into play. Let $V = \{v_0, \dots, v_{n-1}\}$, where the indices are consistent with a breadth-first numbering produced by a breadth-first search starting at node v_0 [179]. The idea is to use the numbering to define an enumeration order: In order to avoid duplicates, the algorithm enumerates connected subgraphs for every node v_i , but restricts them to contain no v_j with $j < i$. Using the definition $\mathcal{B}_i = \{v_j | j \leq i\}$, the pseudocode looks as follows:

EnumerateCsg

Input: a connected query graph $G = (V, E)$

Precondition: nodes in V are numbered according to a breadth-first search

Output: emits all subsets of V inducing a connected subgraph of G

```

for all  $i \in [n - 1, \dots, 0]$  descending {
  emit  $\{v_i\}$ ;
  EnumerateCsgRec( $G, \{v_i\}, \mathcal{B}_i$ );
}

```

`EnumerateCsgRec(G, S, X)`

$N = \mathcal{N}(S) \setminus X$;

for all $S' \subseteq N, S' \neq \emptyset$, enumerate subsets first {

emit $(S \cup S')$;

}

for all $S' \subseteq N, S' \neq \emptyset$, enumerate subsets first {

`EnumerateCsgRec($G, (S \cup S'), (X \cup N)$)`;

}

Let us consider an example. Figure 3.12 contains a query graph whose nodes are numbered in a breadth-first fashion. The calls to `EnumerateCsgRec` are contained in the table in Figure 3.13. In this table, S and X are the arguments of `EnumerateCsgRec`. N is the local variable after its initialization. The column `emit/ S` contains the connected subset emitted, which then becomes the argument of the recursive call to `EnumerateCsgRec` (labelled by \rightarrow). Since listing all calls is too lengthy, only a subset of the calls is listed.

Generating the connected subsets is an important first step but clearly not sufficient: we have to generate all csg-cmp-pairs. The basic idea to do so is as follows. Algorithm `EnumerateCsg` is used to create the first component S_1

EnumerateCsgRec			
S	X	N	emit/ S
{4}	{0, 1, 2, 3, 4}	\emptyset	
{3}	{0, 1, 2, 3}	{4}	{3, 4}
{2}	{0, 1, 2}	{3, 4}	{2, 3}
			{2, 4}
			{2, 3, 4}
{1}	{0, 1}	{4}	{1, 4}
→ {1, 4}	{0, 1, 4}	{2, 3}	{1, 2, 4}
			{1, 3, 4}
			{1, 2, 3, 4}
{0}	{0}	{1, 2, 3}	{0, 1}
			{0, 2}
			{0, 3}
			{0, 1, 2}
			{0, 1, 3}
			{0, 2, 3}
			{0, 1, 2, 3}
→ {0, 1}	{0, 1, 2, 3}	{4}	{0, 1, 4}
→ {0, 2}	{0, 1, 2, 3}	{4}	{0, 2, 4}

Figure 3.13: Call sequence for Figure 3.12

of every csg-cmp-pair. Then, for each such S_1 , we generate all its complement components S_2 . This can be done by calling `EnumerateCsgRec` with the correct parameters. Remember that we have to generate every csg-cmp-pair once and only once.

To achieve this, we use a similar technique as for connected subsets, using the breadth-first numbering to define an enumeration order: we consider only sets S_2 in the complement of S_1 (with (S_1, S_2) being a csg-cmp-pair) such that S_2 contains only v_j with j larger than any i with $v_i \in S_1$. This avoids the generation of duplicates.

We need some definitions to state the actual algorithm. Let $S_1 \subseteq V$ be a non-empty subset of V . Then, we define $\min(S_1) := \min(\{i | v_i \in S_1\})$. This is used to extract the starting node from which S_1 was constructed (see Lemma ??). Let $W \subset V$ be a non-empty subset of V . Then, we define $\mathcal{B}_i(W) := \{v_j | v_j \in W, j < i\}$. Using this notation, the algorithm to construct all S_2 for a given S_1 such that (S_1, S_2) is a csg-cmp-pair looks as follows:

EnumerateCmp

Input: a connected query graph $G = (V, E)$, a connected subset S_1

Precondition: nodes in V are numbered according to a breadth-first search

Output: emits all complements S_2 for S_1 such that (S_1, S_2) is a csg-cmp-pair
 $X = \mathcal{B}_{\min(S_1)} \cup S_1$;

$N = \mathcal{N}(S_1) \setminus X$;

for all ($v_i \in N$ by descending i) {

emit $\{v_i\}$;

 EnumerateCsgRec($G, \{v_i\}, X \cup (\mathcal{B}_i \cap N)$);

}

Algorithm **EnumerateCmp** considers all neighbors of S_1 . First, they are used to determine those S_2 that contain only a single node. Then, for each neighbor of S_1 , it recursively calls **EnumerateCsgRec** to create those S_2 that contain more than a single node. Note that here both nodes concerning the enumeration of S_1 ($\mathcal{B}_{\min(S_1)} \cup S_1$) and nodes concerning the enumeration of S_2 (N) have to be considered in order to guarantee a correct enumeration. Otherwise the combined algorithm would emit (commutative) duplicates.

Let us consider an example for algorithm **EnumerateCmp**. The underlying graph is again the one shown in Fig. 3.12. Assume **EnumerateCmp** is called with $S_1 = \{R_1\}$. In the first statement, the set $\{R_0, R_1\}$ is assigned to X . Then, the neighborhood is calculated. This results in

$$N = \{R_0, R_4\} \setminus \{R_0, R_1\} = \{R_4\}.$$

Hence, $\{R_4\}$ is emitted and together with $\{R_1\}$, it forms the csg-cmp-pair $(\{R_1\}, \{R_4\})$. Then, the recursive call to **EnumerateCsgRec** follows with arguments $G, \{R_4\}$, and $\{R_0, R_1, R_4\}$. Subsequent **EnumerateCsgRec** generates the connected sets $\{R_2, R_4\}$, $\{R_3, R_4\}$, and $\{R_2, R_3, R_4\}$, giving three more csg-cmp-pairs.

3.2.5 Memoization

Whereas dynamic programming constructs the join trees iteratively from small trees to larger trees, i.e. works bottom up, memoization works recursively. For a given set of relations S , it produces the best join tree for S by recursively calling itself for every subset S_1 of S and considering all join trees between S_1 and its complement S_2 . The best alternative is memoized (hence the name). The reason is that two (even different) (sub-) sets of all relations may very well have the common subsets. For example, $\{R_1, R_2, R_3, R_4, R_5\}$ and $\{R_2, R_3, R_4, R_5, R_6\}$ have the common subset $\{R_2, R_3, R_4, R_5\}$. In order to avoid duplicate work, memoization is essential.

In the following variant of memoization, we explore the search space of all bushy trees and consider cross products. We split the functionality across two functions. The first one initializes the **BestTree** data structure with single relation join trees for R_i and then calls the second one. The second one is the core memoization procedure which calls itself recursively.

EX

```

MemoizationJoinOrdering( $R$ )
Input: a set of relations  $R$ 
Output: an optimal join tree for  $R$ 
for ( $i = 1; i \leq n; ++i$ ) {
    BestTree( $\{R_i\}$ ) =  $R_i$ ;
}
return MemoizationJoinOrderingSub( $R$ );

```

```

MemoizationJoinOrderingSub( $S$ )
Input: a (sub-) set of relations  $S$ 
Output: an optimal join tree for  $S$ 
if(NULL == BestTree( $S$ )) {
    for all  $S_1 \subset S$  do {
         $S_2 = S \setminus S_1$ ;
        CurrTree = CreateJoinTree(MemoizationJoinOrderingSub( $S_1$ ), MemoizationJoinOrderingSub( $S_2$ ));
        if (BestTree( $S$ ) == NULL || cost(BestTree( $S$ )) > cost(CurrTree)) {
            BestTree( $S$ ) = CurrTree;
        }
    }
}
return BestTree( $S$ );

```

Again, pruning techniques can help to speed up plan generation [688].

ToDo?

3.2.6 Join Ordering by Generating Permutations

For any set of cost functions, we can directly generate permutations. Generating all permutations is clearly too expensive for more than a couple of relations. However, we can safely neglect some of them. Consider the join sequence $R_1R_2R_3R_4$. If we know that $R_1R_3R_2$ is cheaper than $R_1R_2R_3$, we do not have to consider $R_1R_2R_3R_4$. The idea of the following algorithm is to construct permutations by successively adding relations. Thereby, an extended sequence is only explored if exchanging the last two relations does not result in a cheaper sequence.

```

ConstructPermutations(Query Specification)
Input: query specification for relations  $\{R_1, \dots, R_n\}$ 
Output: optimal left-deep tree
BestPermutation = NULL;
Prefix =  $\epsilon$ ;
Rest =  $\{R_1, \dots, R_n\}$ ;
ConstructPermutationsSub(Prefix, Rest);
return BestPermutation

```



```

ConstructPermutationsSub(Prefix, Rest)
Input: a prefix of a permutation and the relations to be added (Rest)
Output: none, side-effect on BestPermutation
if (Rest ==  $\emptyset$ ) {
    if (BestPermutation == NULL || cost(Prefix) < cost(BestPermutation)) {
        BestPermutation = Prefix;
    }
    return
}
foreach ( $R_i, R_j \in$  Rest) {
    if (cost(Prefix  $\circ$   $\langle R_i, R_j \rangle$ )  $\leq$  cost(Prefix  $\circ$   $\langle R_j, R_i \rangle$ )) {
        ConstructPermutationsSub(Prefix  $\circ$   $\langle R_i \rangle$ , Rest  $\setminus$  { $R_i$ });
    }
    if (cost(Prefix  $\circ$   $\langle R_j, R_i \rangle$ )  $\leq$  cost(Prefix  $\circ$   $\langle R_i, R_j \rangle$ )) {
        ConstructPermutationsSub(Prefix  $\circ$   $\langle R_j \rangle$ , Rest  $\setminus$  { $R_j$ });
    }
}
return

```

The algorithm can be made more efficient, if the `foreach` loop considers only a single relation and performs the swap test with this relation and the last relation occurring in `Prefix`.

The algorithm has two main advantages over dynamic programming and memoization. The first advantage is that it needs only linear space opposed to exponential space for the two mentioned alternatives. The other main advantage over dynamic programming is that it generates join trees early, whereas with dynamic programming we only generate a plan after the whole search space has been explored. Thus, if the query contains too many joins—that is, the search space cannot be fully explored in reasonable time and space—dynamic programming will not generate any plan at all. If stopped, `ConstructPermutations` will not necessarily compute the best plan, but still some plans have been investigated. This allows us to stop it after some time limit has exceeded. The time limit itself can be fixed, like 100 ms, or variable, like 5% of the execution time of the best plan found so far.

The predicates in the `if` statement can be made more efficient if a (local) ranking function is available. Further speed-up of the algorithm can be achieved if additionally the idea of memoization is applied (of course, this jeopardizes the small memory footprint).

The following variant might be interesting if one is willing to go from linear space consumption to quadratic space consumption. The original algorithm is then started n times, once for each relation as a starting relation. The n different instantiations then have to run interleaved. This variant reduces the dependency on the starting relation.

ToDo/EX

Worst Case Analysis

ToDo/EX

Pruning/memoization/propagation

3.2.7 A Dynamic Programming based Heuristics for Chain Queries

In Section 3.1.6, we saw that the complexity of producing optimal left-deep trees possibly containing cross products for chain queries is an open problem. However, the case does not seem to be hopeless. In fact, Scheufele and Morkotte present two algorithms [655, 657] for this problem. For one algorithm, it can be proven that it has polynomial runtime, for the other, it can be proven that it produces the optimal join tree. However, for none of them both could be proven so far.

Basic Definitions and Lemmata

An instance of the *join-ordering problem for chain queries* (or a *chain query* for short) is fully described by the following parameters. First, n relations R_1, \dots, R_n are given. The size of relation R_i ($1 \leq i \leq n$) is denoted by $|R_i|$ or n_{R_i} . Second, the query graph G on the set of relations R_1, \dots, R_n must be a chain. That is, its edges are $\{(R_i, R_{i+1}) \mid 1 \leq i < n\}$:

$$R_1 - R_2 - \dots - R_n$$

For every edge (R_i, R_{i+1}) , there is an associated selectivity $f_{i,i+1} = |R_i \bowtie R_{i+1}| / |R_i \times R_{i+1}|$. We define all other selectivities $f_{i,j} = 1$ for $|i - j| \neq 1$. They correspond to cross products.

In this section we consider only left-deep processing trees. However, we allow them to contain cross products. Hence, any permutation is a valid join tree. There is a unique correspondence not only between left-deep join trees but also between consecutive parts of a permutation and segments of a left-deep tree. Furthermore, if a segment of a left-deep tree does not contain cross products, it uniquely corresponds to a consecutive part of the chain in the query graph. In this case, we also speak of (sub)chains or connected (sub)sequences. We say that two relations R_i and R_j are *connected* if they are adjacent in G ; more generally, two sequences s and t are connected if there exist relations R_i in s and R_j in t such that R_i and R_j are connected. A sequence of relations s is connected if for all subsequences s_1 and s_2 satisfying $s = s_1 s_2$ it holds that s_1 is connected to s_2 .

Given a chain query, we ask for a permutation $s = r_1 \dots r_n$ of the n relations (i.e. there is a permutation π such that $r_i = R_{\pi(i)}$ for $1 \leq i \leq n$) that produces minimal costs under the cost function C_{out} .

Remember that the dynamic programming approach considers $n2^{n-1} - n(n+1)/2$ alternatives for left-deep processing trees with cross products—independently of the query graph and the cost function. The question arises whether it is possible to lower the complexity in case of simple chain queries.

The IKKBZ algorithm solves the join ordering problem for tree queries by decomposing the problem into polynomially many subproblems which are subject to tree-like precedence constraints. The precedence constraints ensure that the cost functions of the subproblems now have the ASI property. The remaining problem is to optimize the constrained subproblems under the simpler cost function. Unfortunately, this approach does not work in our case, since no such decomposition seems to exist.

Let us introduce some notions used for the algorithms. We have to generalize the rank used in the IKKBZ algorithm to *relativized ranks*. We start by relativizing the cost function. The costs of a sequence s *relative* to a sequence u are defined as

$$\begin{aligned} C_u(\epsilon) &:= 0 \\ C_u(R_i) &:= 0 \text{ if } u = \epsilon \\ C_u(R_i) &:= \left(\prod_{R_j <_u R_i} f_{j,i} \right) n_i \text{ if } u \neq \epsilon \\ C_u(s_1 s_2) &:= C_u(s_1) + T_u(s_1) * C_{us_1}(s_2) \end{aligned}$$

with

$$\begin{aligned} T_u(\epsilon) &:= 1 \\ T_u(s) &:= \prod_{R_i \in s} \left(\prod_{R_j <_{us} R_i} f_{j,i} \right) * n_i \end{aligned}$$

Here, $R_i <_s R_j$ is true if and only if R_i appears before R_j in s . As usual, empty products evaluate to 1. Several things should be noted. First, $C_{us}(t) = C_u(t)$ holds if there is no connection between relations in s and t . Second, $T_\epsilon(R_i) = |R_i|$ and $T_\epsilon(s) = |s|$. That is, T_u generalizes the size of a single relation or of a sequence of relations. Third, note that $C_u(\epsilon) = 0$ for all u but $C_\epsilon(s) = 0$ only if s does not contain more than one relation. The special case that $C_\epsilon(R) = 0$ for a single relation R causes some problems in the homogeneity of definitions and proofs. Hence, we abandon this case from all definitions and lemmata of this section. This will not be repeated in every definition and lemma, but will implicitly be assumed. Further, the two algorithms will be presented in two versions. The first version is simpler and relies on a modified cost function C' , and only the second version will apply to the original cost function C . As we will see, C' differs from C in exactly the problematic case in which it is defined as $C'_u(R_i) := |R_i|$. Now, $C'_\epsilon(s) = 0$ holds if and only if $s = \epsilon$ holds. Within subsequent definitions and lemmata, C can also be replaced by C' without changing their validity. Last, we abbreviate C_ϵ by C for convenience.

Example 1: Consider a chain query involving the relations R_1, R_2, R_3 . The parameters are $|R_1| = 1, |R_2| = 100, |R_3| = 10$ and $f_{1,2} = f_{2,3} = 0.9$. The expected size of the query result is independent of the ordering of the relations. Hence, we have

$$T(R_1 R_2 R_3) = \dots = T(R_3 R_2 R_1) = 100 * 10 * 1 * .9 * .9 = 810.$$

There are 6 possible orderings of the relations with the following costs:

$$\begin{aligned} C(R_1R_2R_3) &= 1 * 100 * 0.9 + 1 * 100 * 10 * 0.9 * 0.9 = 900 \\ C(R_1R_3R_2) &= 1 * 10 + 1 * 10 * 100 * 0.9 * 0.9 = 820 \\ C(R_2R_3R_1) &= 100 * 10 * 0.9 + 100 * 10 * 1 * 0.9 * 0.9 = 1710 \\ C(R_2R_1R_3) &= C(R_1R_2R_3) \\ C(R_3R_1R_2) &= C(R_1R_3R_2) \\ C(R_3R_2R_1) &= C(R_2R_3R_1) \end{aligned}$$

Note that the cost function is invariant with respect to the order of the first two relations. The minimum over all costs is 820, and the corresponding optimal join ordering is $R_1R_3R_2$. □

Using the relativized cost function, we can define the relativized rank.

Definition 3.2.8 (rank) *The rank of a sequence s relative to a non-empty sequence u is given by*

$$rank_u(s) := \frac{T_u(s) - 1}{C_u(s)}$$

In the special case that s consists of a single relation R_i , the intuition behind the *rank* function becomes transparent. Let f_i be the product of the selectivities between relations in u and R_i . Then $rank_u(R_i) = \frac{f_i|R_i|-1}{f_i|R_i|}$. Hence, the *rank* becomes a function of the form $f(x) = \frac{x-1}{x}$. This function is monotonously increasing in x for $x > 0$. The argument to the function $f(x)$ is (for the computation of the size of a single relation R_i) $f_i|R_i|$. But this is the factor by which the next intermediate result will increase (or decrease). Since we sum up intermediate results, this is an essential number. Furthermore, it follows from the monotonicity of $f(x)$ that $rank_u(R_i) \leq rank_u(R_j)$ if and only if $f_i|R_i| \leq f_j|R_j|$ where f_j is the product of all selectivities between R_j and relations in u .

Example 1 (cont'd): Supposing the query given in Example 1, the optimal sequence $R_1R_3R_2$ gives rise to the following ranks.

$$\begin{aligned} rank_{R_1}(R_2) &= \frac{T_{R_1}(R_2)-1}{C_{R_1}(R_2)} = \frac{100*0.9-1}{100*0.9} \approx 0.9888 \\ rank_{R_1}(R_3) &= \frac{T_{R_1}(R_3)-1}{C_{R_1}(R_3)} = \frac{10*1.0-1}{10*1.0} = 0.9 \\ rank_{R_1R_3}(R_2) &= \frac{T_{R_1R_3}(R_2)-1}{C_{R_1R_3}(R_2)} = \frac{100*0.9*0.9-1}{100*0.9*0.9} \approx 0.9877 \end{aligned}$$

Hence, within the optimal sequence, the relation with the smallest rank (here R_3 , since $rank_{R_1}(R_3) < rank_{R_1}(R_2)$) is preferred. As the next lemma will show, this is no accident. □

Using the rank function, the following lemma can be proved.

Lemma 3.2.9 *For sequences*

$$\begin{aligned} S &= r_1 \cdots r_{k-1} r_k r_{k+1} r_{k+2} \cdots r_n \\ S' &= r_1 \cdots r_{k-1} r_{k+1} r_k r_{k+2} \cdots r_n \end{aligned}$$

the following holds:

$$C(S) \leq C(S') \Leftrightarrow \text{rank}_u(r_k) \leq \text{rank}_u(r_{k+1})$$

where $u = r_1 \cdots r_{k-1}$. Equality only holds if it holds on both sides.

Example 1 (cont'd): Since the ranks of the relations in Example 1 are ordered with ascending ranks, Lemma 3.2.9 states that, whenever we exchange two adjacent relations, the costs cannot decrease. In fact, we observe that $C(R_1R_3R_2) \leq C(R_1R_2R_3)$. \square

An analogous lemma still holds for two unconnected subchains:

Lemma 3.2.10 *Let u, x and y be three subchains where x and y are not interconnected. Then we have:*

$$C(uxy) \leq C(uyx) \Leftrightarrow \text{rank}_u(x) \leq \text{rank}_u(y)$$

Equality only holds if it holds on both sides.

Next, we define the notion of a *contradictory chain*, which will be essential to the algorithms. The subsequent lemmata will allow us to cut down the search space to be explored by any optimization algorithm.

Definition 3.2.11 (contradictory pair of subchains) *Let u, x, y be nonempty sequences. We call (x, y) a contradictory pair of subchains if and only if*

$$C_u(xy) \leq C_u(yx) \quad \wedge \quad \text{rank}_u(x) > \text{rank}_{ux}(y)$$

A special case occurs when x and y are single relations. Then the above condition simplifies to

$$\text{rank}_{ux}(y) < \text{rank}_u(x) \leq \text{rank}_u(y)$$

To explain the intuition behind the definition of contradictory subchains, we need another example.

Example 2: Suppose a chain query involving R_1, R_2, R_3 is given. The relation sizes are $|R_1| = 1, |R_2| = |R_3| = 10$ and the selectivities are $f_{1,2} = 0.5, f_{2,3} = 0.2$. Consider the sequences $R_1R_2R_3$ and $R_1R_3R_2$, which differ in the order of the last two relations. We have

$$\begin{aligned} \text{rank}_{R_1}(R_2) &= 0.8 \\ \text{rank}_{R_1R_2}(R_3) &= 0.0 \\ \text{rank}_{R_1}(R_3) &= 0.9 \\ \text{rank}_{R_1R_3}(R_2) &= 0.5 \end{aligned}$$

and

$$\begin{aligned} C(R_1R_2R_3) &= 15 \\ C(R_1R_3R_2) &= 20 \end{aligned}$$

Hence,

$$\begin{aligned} \text{rank}_{R_1}(R_2) &> \text{rank}_{R_1R_2}(R_3) \\ \text{rank}_{R_1}(R_3) &> \text{rank}_{R_1R_3}(R_2) \\ C(R_1R_2R_3) &< C(R_1R_3R_2) \end{aligned}$$

and (R_2, R_3) is a contradictory pair within $R_1R_2R_3$. Now the use of the term *contradictory* becomes clear: the costs do not behave as could be expected from the ranks. \square

The next (obvious) lemma states that contradictory chains are necessarily connected.

Lemma 3.2.12 *If there is no connection between two subchains x and y , then they cannot build a contradictory pair (x, y) .*

Now we present the fact that between a contradictory pair of relations, there cannot be any other relation not connected to them without increasing cost.

Lemma 3.2.13 *Let $S = usvtw$ be a sequence. If there is no connection between relations in s and v and relations in v and t , and $\text{rank}_u(s) \geq \text{rank}_{us}(t)$, then there exists a sequence S' not having higher costs, where s immediately precedes t .*

Example 3: Consider five relations R_1, \dots, R_5 . The relation sizes are $|R_1| = 1$, $|R_2| = |R_3| = |R_4| = 8$, and $|R_5| = 2$. The selectivities are $f_{1,2} = \frac{1}{2}$, $f_{2,3} = \frac{1}{4}$, $f_{3,4} = \frac{1}{8}$, and $f_{4,5} = \frac{1}{2}$. Relation R_5 is not connected to relations R_2 and R_3 . Further, within the sequence $R_1R_2R_5R_3R_4$ relations R_2 and R_3 have contradictory ranks: $\text{rank}_{R_1}(R_2) = \frac{4-1}{4} = \frac{3}{4}$ and $\text{rank}_{R_1R_2R_5}(R_3) = \frac{2-1}{2} = \frac{1}{2}$. Hence, at least one of $R_1R_5R_2R_3R_4$ and $R_1R_2R_3R_5R_4$ must be of no greater cost than $R_1R_2R_5R_3R_4$. This is indeed the case:

$$\begin{aligned} C(R_1R_2R_3R_5R_4) &= 4 + 8 + 16 + 8 = 36 \\ C(R_1R_2R_5R_3R_4) &= 4 + 8 + 16 + 8 = 36 \\ C(R_1R_5R_2R_3R_4) &= 2 + 8 + 16 + 8 = 34 \end{aligned}$$

\square

The next lemma shows that, if there exist two sequences of single rank-sorted relations, then their costs as well as their ranks are necessarily equal.

Lemma 3.2.14 *Let $S = x_1 \dots x_n$ and $S' = y_1 \dots y_n$ be two different rank-sorted chains containing exactly the relations R_1, \dots, R_n , i.e.*

$$\begin{aligned} \text{rank}_{x_1 \dots x_{i-1}}(x_i) &\leq \text{rank}_{x_1 \dots x_i}(x_{i+1}) \quad \text{for all } 1 \leq i \leq n, \\ \text{rank}_{y_1 \dots y_{i-1}}(y_i) &\leq \text{rank}_{y_1 \dots y_i}(y_{i+1}) \quad \text{for all } 1 \leq i \leq n, \end{aligned}$$

then S and S' have equal costs and, furthermore,

$$\text{rank}_{x_1 \dots x_{i-1}}(x_i) = \text{rank}_{y_1 \dots y_{i-1}}(y_i) \quad \text{for all } 1 < i \leq n$$

One could conjecture that the following generalization of Lemma 3.2.14 is true, although no one has proved it so far.

Conjecture 3.2.1 *Let $S = x_1 \cdots x_n$ and $S' = y_1 \cdots y_m$ be two different rank-sorted chains for the relations $R_1 \dots, R_n$ where the x'_i 's and y'_i 's are subsequences such that*

$$\begin{aligned} \text{rank}_{x_1 \dots x_{i-1}}(x_i) &\leq \text{rank}_{x_1 \dots x_i}(x_{i+1}) \text{ for all } 1 \leq i < n, \\ \text{rank}_{y_1 \dots y_{i-1}}(y_i) &\leq \text{rank}_{y_1 \dots y_i}(y_{i+1}) \text{ for all } 1 \leq i < m, \end{aligned}$$

and the subsequences x_i and y_j are all optimal (with respect to the fixed prefixes $x_1 \dots x_{i-1}$ and $y_1 \dots y_{j-1}$), then S and S' have equal costs.

Consider the problem of merging two optimal unconnected chains. If we knew that the ranks of relations in an optimal chain are always sorted in ascending order, we could use the classical merge procedure to combine the two chains. The resulting chain would also be rank-sorted in ascending order and, according to Lemma 3.2.14, it would be optimal. Unfortunately, this does not work, since there are optimal chains whose ranks are not sorted in ascending order: those containing sequences with contradictory ranks.

Now, as shown in Lemma 3.2.13, between contradictory pairs of relations there cannot be any other relation not connected to them. Hence, in the merging process, we have to take care that we do not merge a contradictory pair of relations with a relation not connected to the pair. In order to achieve this, we apply the same trick as in the IKKBZ algorithm: we tie the relations of a contradictory subchain together by building a *compound relation*. Assume that we tie together the relations r_1, \dots, r_n to a new relation $r_{1, \dots, n}$. Then we define the size of $r_{1, \dots, n}$ as $|r_{1, \dots, n}| = |r_1 \bowtie \dots \bowtie r_n|$. Further, if some r_i ($1 \leq i \leq n$) does have a connection to some $r_k \notin \{r_1, \dots, r_n\}$ then we define the selectivity factor $f_{r_{1, \dots, n}, r_k}$ between r_k and $r_{1, \dots, n}$ as $f_{r_{1, \dots, n}, r_k} = f_{i, k}$.

If we tie together contradictory pairs, the resulting chain of compound relations still does not have to be rank-sorted with respect to the compound relations. To overcome this, we iterate the process of tying contradictory pairs of compound relations together until the sequence of compound relations is rank-sorted, which will eventually be the case. That is, we apply the *normalization* as used in the IKKBZ algorithm. However, we have to reformulate it for relativized costs and ranks:

Normalize(p, s)

```

while (there exist subsequences  $u, v$  ( $u \neq \epsilon$ ) and
        compound relations  $x, y$  such that  $s = uxyv$ 
        and  $C_{pu}(xy) \leq C_{pu}(yx)$ 
        and  $\text{rank}_{pu}(x) > \text{rank}_{pux}(y)$ ) {
    replace  $xy$  by a compound relation  $(x, y)$ ;
}
return ( $p, s$ );

```

The compound relations in the result of the procedure `Normalize` are called *contradictory chains*. A *maximal contradictory subchain* is a contradictory subchain that cannot be made longer by further tying steps. Resolving the tyings introduced in the procedure `normalize` is called *de-normalization*. It works the same way as in the IKKBZ algorithm. The cost, size and rank functions can now be extended to sequences containing compound relations in a straightforward way. We define the cost of a sequence containing compound relations to be identical with the cost of the corresponding de-normalized sequence. The size and rank functions are defined analogously.

The following simple observation is central to the algorithms: every chain can be decomposed into a sequence of adjacent maximal contradictory subchains. For convenience, we often speak of chains instead of subchains and of contradictory chains instead of maximal contradictory subchains. The meaning should be clear from the context. Further, we note that the decomposition into adjacent maximal contradictory subchains is not unique. For example, consider an optimal subchain $r_1 r_2 r_3$ and a sequence u of preceding relations. If $rank_u(r_1) > rank_{ur_1}(r_2) > rank_{ur_1 r_2}(r_3)$ one can easily show that both $(r_1, (r_2, r_3))$ and $((r_1, r_2), r_3)$ are contradictory subchains. Nevertheless, this ambiguity is not important since in the following we are only interested in contradictory subchains which are *optimal*. In this case, the condition $C_u(xy) \leq C_u(yx)$ is certainly true and can therefore be neglected. One can show that for the case of optimal subchains the indeterministically defined normalization process is well-defined, that is, if S is optimal, `normalize(P, S)` will always terminate with a unique “flat” decomposition of S into maximal contradictory subchains (flat means that we remove all but the outermost parenthesis, e.g. $(R_1 R_2)((R_5 R_4) R_3) R_6$) becomes $(R_1 R_2)(R_5 R_4 R_3 R_6)$).

The next two lemmata and the conjecture show a possible way to overcome the problem that if we consider cross products, we have an unconstrained ordering problem and the idea of Monma and Sidney as exploited in the IKKBZ algorithm is no longer applicable. The next lemma is a direct consequence of the normalization procedure.

Lemma 3.2.15 *Let $S = s_1 \dots s_m$ be an optimal chain consisting of the maximal contradictory subchains s_1, \dots, s_m (as determined by the function `normalize`). Then*

$$\begin{aligned} rank(s_1) &\leq rank_{s_1}(s_2) \leq rank_{s_1 s_2}(s_3) \\ &\leq \dots \leq rank_{s_1 \dots s_{m-1}}(s_m), \end{aligned}$$

in other words, the (maximal) contradictory subchains in an optimal chain are always sorted by ascending ranks.

The next result shows how to build an optimal sequence from two optimal non-interconnected sequences.

Lemma 3.2.16 *Let x and y be two optimal sequences of relations where x and y are not interconnected. Then the sequence obtained by merging the maximal contradictory subchains in x and y (as obtained by `normalize`) according to their ascending rank is optimal.*

Merging two sequences in the way described in Lemma 3.2.16 is a fundamental process. We henceforth refer to it by simply saying that we *merge by the ranks*.

We strongly conjecture that the following generalization of Lemma 3.2.14 is true, although it is yet unproven. It uses the notion of *optimal recursive decomposable subchains* defined in the next subsection.

Conjecture 3.2.2 *Consider two sequences S and T containing exactly the relations R_1, \dots, R_n . Let $S = s_1 \dots s_k$ and $T = t_1 \dots t_l$ be such that each of the maximal contradictory subchains $s_i, i = 1, \dots, k$ and $t_j, j = 1, \dots, l$ are optimal recursively decomposable. Then S and T have equal costs.*

The first algorithm

We first use a slightly modified cost function C' , which additionally respects the size of the first relation in the sequence, i.e. C and C' relate via

$$C'_u(s) = \begin{cases} C(s) + |n_R|, & \text{if } u = \epsilon \text{ and } s = Rs' \\ C_u(s), & \text{otherwise} \end{cases}$$

This cost function can be treated in a more elegant way than C . The new rank function is now defined as $rank_u(s) := (T_u(s) - 1)/C'_u(s)$. Note that the rank function is now defined even if $u = \epsilon$ and s is a single relation. The size function remains unchanged. At the end of this subsection, we describe how our results can be adapted to the original cost function C .

The rank of a contradictory chain depends on the relative position of the relations that are directly connected to it. For example, the rank of the contradictory subchain $(R_5R_3R_4R_2)$ depends on the position of the neighbouring relations R_1 and R_6 relative to $(R_5R_3R_4R_2)$. That is, whether they appear before or after the sequence $(R_5R_3R_4R_2)$. Therefore, we introduce the following fundamental definitions:

Definition 3.2.17 (neighbourhood) *We call the set of relations that are directly connected to a subchain (with respect to the query graph G) the complete neighbourhood of that subchain. A neighbourhood is a subset of the complete neighbourhood. The complement of a neighbourhood u of a subchain s is defined as $v \setminus u$, where v denotes the complete neighbourhood of s .*

Note that the neighbourhood of a subchain s within a larger chain us is uniquely determined by the subsequence u of relations preceding it. For convenience, we will often use sequences of preceding relations to specify neighbourhoods. We henceforth denote a pair consisting of a connected sequence s and a neighbourhood u by $[s]_u$.

Definition 3.2.18 (contradictory subchain, extent) *A contradictory subchain $[s]_u$ is inductively defined as follows.*

1. For a single relation s , $[s]_u$ is a contradictory subchain.

2. There is a decomposition $s = vw$ such that (v, w) is a contradictory pair with respect to the preceding subsequence u and both $[v]_u$ and $[w]_{uv}$ are contradictory subchains themselves.

The extent of a contradictory chain $[s]_u$ is defined as the pair consisting of the neighbourhood u and the set of relations occurring in s . Since contradictory subchains are connected, the set of occurring relations has always the form $\{R_i, R_{i+1}, \dots, R_{i+l}\}$ for some $1 \leq i \leq n$, $0 \leq l \leq n - i$. An optimal contradictory subchain to a given extent is a contradictory subchain with lowest cost among all contradictory subchains of the same extent.

The number of different extents of contradictory subchains for a chain query of n relations is $2n^2 - 2n + 1$. Each contradictory chain can be completely recursively decomposed into adjacent pairs of connected subchains. Subchains with this property are defined next (similar types of decompositions occur in [374, 689]).

Definition 3.2.19 ((optimal) recursively decomposable subchain) A recursively decomposable subchain $[s]_u$ is inductively defined as follows.

1. If s is a single relation, then $[s]_u$ is recursively decomposable.
2. There is a decomposition $s = vw$ such that v is connected to w and both $[v]_u$ and $[w]_{uv}$ are recursively decomposable subchains.

The extent of a recursively decomposable chain is defined in the same way as for contradictory chains. Note that every contradictory subchain is recursively decomposable. Consequently, the set of all contradictory subchains for a certain extent is a subset of all recursively decomposable subchains of the same extent.

Example 4: Consider the sequence of relations

$$s = R_2R_4R_3R_6R_5R_1.$$

Using parentheses to indicate the recursive decompositions, we have the following two possibilities

$$(((R_2(R_4R_3))(R_6R_5))R_1)$$

$$((R_2((R_4R_3)(R_6R_5)))R_1)$$

The extent of the recursively decomposable subchain $R_4R_3R_6R_5$ of s is $(\{R_2\}, \{R_3, R_4, R_5, R_6\})$. \square

The number of different recursively decomposable chains involving the relations R_1, \dots, R_n is r_n , where r_n denotes the n -th Schröder number [689]. Hence, the total number of recursively decomposable chains is $r_n + 2(n - 1)r_{n-1} + 4 \sum_{i=1}^{n-2} \binom{n-2}{i} r_i$. It can be shown that

$$r_n \approx \frac{C(2 + \sqrt{8})^n}{n^{3/2}}$$

where $C = 1/2\sqrt{\frac{2\sqrt{2}-4}{\pi}}$. Using Stirling's formula for $n!$ it is easy to show that $\lim_{n \rightarrow \infty} \frac{r_n}{n!} = 0$. Thus, the probability of a random permutation to be recursively decomposable strives to zero for large n .

An *optimal recursively decomposable subchain* to a given extent is a recursively decomposable subchain with lowest cost among all recursively decomposable subchains of the same extent. There is an obvious dynamic programming algorithm to compute optimal recursive decomposable subchains. It is not hard to see that *Bellman's optimality principle* [520, 179] holds and every optimal recursively decomposable subchain can be decomposed into smaller optimal recursively decomposable subchains.

Example 5: In order to compute an optimal recursively decomposable subchain for the extent

$$(\{R_2, R_7\}, \{R_3, R_4, R_5, R_6\})$$

the algorithm makes use of optimal recursively decomposable subchains for the extents

$$\begin{array}{ll} (\{R_2\}, \{R_3\}) & (\{R_7, R_3\}, \{R_4, R_5, R_6\}) \\ (\{R_2\}, \{R_3, R_4\}) & (\{R_7, R_4\}, \{R_5, R_6\}) \\ (\{R_2\}, \{R_3, R_4, R_5\}) & (\{R_5, R_7\}, \{R_6\}) \\ (\{R_7\}, \{R_4, R_5, R_6\}) & (\{R_2, R_4\}, \{R_3\}) \\ (\{R_7\}, \{R_5, R_6\}) & (\{R_2, R_5\}, \{R_3, R_4\}) \\ (\{R_7\}, \{R_6\}) & (\{R_2, R_6\}, \{R_3, R_4, R_5\}) \end{array}$$

which have been computed in earlier steps². A similar dynamic programming algorithm can be used to determine optimal contradictory subchains. \square

Let E be the set of all possible extents. We define the following partial order $\mathcal{P} = (E, \prec)$ on E . For all extents $e_1, e_2 \in E$, we have $e_1 \prec e_2$ if and only if e_1 can be obtained by splitting the extent e_2 . For example, $(\{R_7\}, \{R_5, R_6\}) \prec (\{R_2, R_7\}, \{R_3, R_4, R_5, R_6\})$. The set of maximal extents M then corresponds to a set of incomparable elements (antichain) in \mathcal{P} such that for all extents e enumerated so far, there is an extent $e' \in M$ with $e \prec e'$.

Now, since every optimal join sequence has a representation as a sequence of contradictory subchains, we only have to determine this representation. Consider a contradictory subchain c in an optimal join sequence s . What can we say about c ? Obviously, c has to be optimal with respect to the neighbourhood defined by the relations preceding c in s . Unfortunately, identifying contradictory subchains that are optimal sequences seems to be as hard as the whole problem of optimizing chain queries. Therefore, we content ourselves with the following weaker condition which may lead to multiple representations. Nevertheless, it seems to be the strongest condition for which all subchains satisfying the condition can be computed in polynomial time. The condition says that s should be optimal both with respect to all contradictory chains of the same extent as s and with respect to all recursively decomposable subchains of the same extent.

²The splitting of extents induces a partial order on the set of extents.

So far it is not clear whether these conditions lead to multiple representations. Therefore, we have no choice but to enumerate all possible representations and select the one with minimal costs. Next we describe the first algorithm.

Algorithm Chain-I’:

1. Use dynamic programming to determine all optimal contradictory subchains.
This step can be made faster by keeping track of the set M of all maximal extents (with respect to the partial order induced by splitting extents).
2. Determine all optimal recursively decomposable subchains for all extents included in some maximal extent in M .
3. Compare the results from steps 1 and 2 and retain only matching subchains.
4. Sort the contradictory subchains according to their ranks.
5. Eliminate contradictory subchains that cannot be part of a solution.
6. Use backtracking to enumerate all sequences of rank-ordered optimal contradictory subchains and keep track of the sequence with lowest cost.

In step 5 of the algorithm, we eliminate contradictory subchains that do not contribute to a solution. Note that the contradictory subchains in an optimal sequence are characterized by the following two conditions.

1. The extents of all contradictory subchains in the representation build a partition of the set of all relations.
2. The neighbourhoods of all contradictory subchains are consistent with the relations occurring at earlier and later positions in the sequence.

Note that any contradictory subchain occurring in the optimal sequence (except at the first and last positions) necessarily has matching contradictory subchains preceding and succeeding it in the list. In fact, every contradictory subchain X occurring in the optimal join sequence must satisfy the following two conditions.

1. For every relation R in the neighbourhood of X , there exists a contradictory subchain Y at an earlier position in the list which itself meets condition 1, such that R occurs in Y , and Y can be followed by X .
2. For every relation R in the complementary neighbourhood of X , there exists a contradictory subchain Y at a later position in the list which itself meets condition 2, such that R occurs in the neighbourhood of Y , and X can be followed by Y .

Using these two conditions, we can eliminate “useless” contradictory chains from the rank-ordered list by performing a reachability algorithm for each of the DAGs defined by the conditions 1 and 2. In the last step of our algorithm, backtracking is used to enumerate all representations. Suppose that at some step of the algorithm we have determined an initial sequence of contradictory subchains and have a rank-sorted list of the remaining possible contradictory subchains. In addition to the two conditions mentioned above, another reachability algorithm can be applied to determine the set of reachable relations from

the list (with respect to the given prefix). With the use of this information, all branches that do not lead to a complete join sequence can be pruned.

Let us analyze the worst case time complexity of the algorithm. The two dynamic programming steps both iterate over $O(n^2)$ different extents, and each extent gives rise to $O(n)$ splittings. Moreover, for each extent one normalization is necessary, which requires linear time (cost, size and rank can be computed in constant time using recurrences). Therefore, the complexity of the two dynamic programming steps is $O(n^4)$. Sorting $O(n^2)$ contradictory chains can be done in time $O(n^2 \log n)$. The step where all “useless” contradictory subchains are eliminated, consists of two stages of a reachability algorithm which has complexity $O(n^4)$. If conjecture 3.2.2 is true, the backtracking step requires linear time, and the total complexity of the algorithm is $O(n^4)$. Otherwise, if conjecture 3.2.2 is false, the algorithm might exhibit exponential worst case time complexity.

We now describe how to reduce the problem for our original cost function C to the problem for the modified cost function C' . One difficulty with the original cost function is that the ranks are defined only for subsequences of *at least two* relations. Hence, for determining the first relation in our solution we do not have sufficient information. An obvious solution to this problem is to try every relation as starting relation, process each of the two resulting chain queries separately and choose the chain with minimum costs. The new complexity will increase by about a factor of n . This first approach is not very efficient, since the dynamic programming computations overlap considerably, e.g. if we perform dynamic programming on the two overlapping chains $R_1R_2R_3R_4R_5R_6$ and $R_2R_3R_4R_5R_6R_7$, for the intersecting chain $R_2R_3R_4R_5R_6$ everything is computed twice. The cue is that we can perform the dynamic programming calculations before we consider a particular starting relation. Hence, the final algorithm can be sketched as follows:

Algorithm CHAIN-I:

1. Compute all optimal contradictory chains by dynamic programming (corresponds to the steps 1-4 of Algorithm I')
2. For each starting relation R_i , perform the following steps:
 - (a) Let L_1 be the result of applying steps 5 and 6 of Algorithm I' to all contradictory subchains whose extent (N, M) satisfies $R_i \in N$ and $M \subseteq \{R_1, \dots, R_i\}$.
 - (b) Let L_2 be the result of applying steps 5 and 6 of Algorithm I' to all contradictory subchains whose extent (N, M) satisfies $R_i \in N$ and $M \subseteq \{R_i, \dots, R_n\}$.
 - (c) For all $(l_1, l_2) \in L_1 \times L_2$, perform the following steps:
 - i. Let L be the result of merging l_1 and l_2 according to their ranks.
 - ii. Use R_iL to update the current-best join ordering.

Suppose that conjecture 3.2.2 is true, and we can replace the backtracking part by a search for the first solution. Then the complexity of the step 1 is $O(n^4)$,

whereas the complexity of step 2 amounts to $\sum_{i=1}^n (O(i^2) + O(n-i)^2 + O(n)) = O(n^3)$. Hence, the total complexity would be $O(n^4)$ in the worst case. Of course, if our conjecture is false, the necessary backtracking step might lead to an exponential worst case complexity.

The second algorithm

The second algorithm is much simpler than the first one but proves to be less efficient in practice. Since the new algorithm is very similar to some parts of the old one, we just point out the differences between both algorithms. The new version of the algorithm works as follows.

Algorithm CHAIN-II':

1. Use dynamic programming to compute an optimal recursive decomposable chain for the whole set of relations $\{R_1, \dots, R_n\}$.
2. Normalize the resulting chain.
3. Reorder the contradictory subchains according to their ranks.
4. De-normalize the sequence.

Step 1 is identical to step 2 of our first algorithm. Note that Lemma 3.2.15 cannot be applied to the sequence in Step 2, since an optimal recursive decomposable chain is not necessarily an optimal chain. Therefore, the question arises whether Step 3 really makes sense. One can show that the partial order defined by the precedence relation among the contradictory subchains has the property that all elements along paths in the partial order are sorted by rank. By computing a greedy topological ordering (greedy with respect to the ranks), we obtain a sequence as requested in step 3.

Let us briefly analyze the worst case time complexity of the second algorithm. The first step requires time $O(n^4)$, whereas the second step requires time $O(n^2)$. The third step has complexity $O(n \log n)$. Hence, the total complexity is $O(n^4)$.

Algorithm II' is based on the cost function C' . We can now modify the algorithm for the original cost function C as follows.

Algorithm CHAIN-II:

1. Compute all optimal recursive decomposable chains by dynamic programming (corresponds to step 1 of Algorithm II')
2. For each starting relation R_i , perform the following steps:
 - (a) Let L_1 be the result of applying the steps 2 and 3 of Algorithm II' to all optimal recursive decomposable subchains whose extent (N, M) satisfies $R_i \in N$ and $M \subseteq \{R_1, \dots, R_i\}$.
 - (b) Let L_2 be the result of applying the steps 2 and 3 of Algorithm II' to all optimal recursive decomposable subchains whose extent (N, M) satisfies $R_i \in N$ and $M \subseteq \{R_i, \dots, R_n\}$.
 - (c) Let L be the result of merging L_1 and L_2 according to their ranks.

- (d) De-normalize L .
- (e) Use R_iL to update the current-best join ordering.

The complexity of Step 1 is $O(n^4)$, whereas the complexity of Step 2 amounts to $\sum_{i=1}^n (O(i^2) + O(n-i)^2 + O(n)) = O(n^3)$. Hence, the time complexity of Algorithm II is $O(n^4)$.

Summarizing, we are now left with one algorithm that produces the optimal result but whose worst-case runtime behavior is unknown and one algorithm with polynomial runtime but producing a result which has not been proven to be optimal. Due to this lack of hard facts, Moerkotte and Scheufele ran about 700,000 experiments with random queries of sizes up to 30 relations and fewer experiments for random queries with up to 300 relations to compare the results of our algorithms. For $n \leq 15$, they additionally compared the results with a standard dynamic programming algorithm. Their findings can be summarized as follows.

- All algorithms yielded identical results.
- Backtracking always led to exactly one sequence of contradictory chains.
- In the overwhelming majority of cases the first algorithm proved to be faster than the second.

Whereas the run time of the second algorithm is mainly determined by the number of relations in the query, the run time of the first also heavily depends on the number of existing optimal contradictory subchains. In the worst case, the first algorithm is slightly inferior to the second. Additionally, Hamalainen reports on an independent implementation of the second algorithm [338]. He could not find an example where the second algorithm did not produce the optimal result either. We encourage the reader to prove that it produces the optimal result.

EX

3.2.8 Transformation-Based Approaches

The idea of transformation-based algorithms can be described as follows. Starting from an arbitrary join tree, equivalences (such as commutativity and associativity) are applied to it to derive a set of new join trees. For each of the join trees, the equivalences are again applied to derive even more join trees. This procedure is repeated until no new join tree can be derived. This procedure exhaustively enumerates the set of all bushy trees. Furthermore, before an equivalence is applied, it is difficult to see whether the resulting join tree has already been produced or not (see also Figure 2.6). Thus, this procedure is highly inefficient. Hence, it does not play any role in practice. Nevertheless, we give the pseudo-code for it, since it forms the basis for several of the following algorithms. We split the exhaustive transformation approach into two algorithms. One that applies all equivalences to a given join tree (`ApplyTransformations`) and another that does the loop (`ExhaustiveTransformation`). A transformation is applied in a directed way. Thus, we reformulate commutativity and associativity as rewrite rules using \rightsquigarrow to indicate the direction.

The following table summarizes all rules commonly used in transformation-based and randomized join ordering algorithms. The first three are directly derived from the commutativity and associativity laws for the join. The other rules are shortcuts used under special circumstances. For example, left associativity may turn a left-deep tree into a bushy tree. When only left-deep trees are to be considered, we need a replacement for left associativity. This replacement is called left join exchange.

$R_1 \bowtie R_2$	\rightsquigarrow	$R_2 \bowtie R_1$	Commutativity
$(R_1 \bowtie R_2) \bowtie R_3$	\rightsquigarrow	$R_1 \bowtie (R_2 \bowtie R_3)$	Right Associativity
$R_1 \bowtie (R_2 \bowtie R_3)$	\rightsquigarrow	$(R_1 \bowtie R_2) \bowtie R_3$	Left Associativity
$(R_1 \bowtie R_2) \bowtie R_3$	\rightsquigarrow	$(R_1 \bowtie R_3) \bowtie R_2$	Left Join Exchange
$R_1 \bowtie (R_2 \bowtie R_3)$	\rightsquigarrow	$R_2 \bowtie (R_1 \bowtie R_3)$	Right Join Exchange

Two more rules are often used to transform left-deep trees. The first operation (*swap*) exchanges two arbitrary relations in a left-deep tree. The second operation (*3Cycle*) performs a cyclic rotation of three arbitrary relations in a left-deep tree. To account for different join methods, a rule called *join method exchange* is introduced.

The first rule set (*RS-0*) we are using contains the commutativity rule and both associativity rules. Applying associativity can lead to cross products. RS-0 If we do not want to consider cross products, we only apply any of the two associativity rules if the resulting expression does not contain a cross product. It is easy to extend `ApplyTransformations` to cover this by extending the if conditions with

```
and (ConsiderCrossProducts || connected(.))
```

where the argument of `connected` is the result of applying a transformation.

```
ExhaustiveTransformation({R1,...,Rn})
Input: a set of relations
Output: an optimal join tree
Let T be an arbitrary join tree for all relations
Done = {}; // contains all trees processed
ToDo = {T}; // contains all trees to be processed
while (!empty(ToDo)) {
  Let T be an arbitrary tree in ToDo
  ToDo \ = T;
  Done \cup = T;
  Trees = ApplyTransformations(T);
  for all T ∈ Trees do {
    if (T ∉ ToDo ∪ Done) {
      ToDo += T;
    }
  }
}
```



```
return cheapest tree found in Done;
```

```
ApplyTransformations( $T$ )
```

```
Input: join tree
```

```
Output: all trees derivable by associativity and commutativity
```

```
Trees =  $\emptyset$ ;
```

```
Subtrees = all subtrees of  $T$  rooted at inner nodes
```

```
for all  $S \in$  Subtrees do {
  if ( $S$  is of the form  $S_1 \bowtie S_2$ ) {
    Trees +=  $S_2 \bowtie S_1$ ;
  }
  if ( $S$  is of the form  $(S_1 \bowtie S_2) \bowtie S_3$ ) {
    Trees +=  $S_1 \bowtie (S_2 \bowtie S_3)$ ;
  }
  if ( $S$  is of the form  $S_1 \bowtie (S_2 \bowtie S_3)$ ) {
    Trees +=  $(S_1 \bowtie S_2) \bowtie S_3$ ;
  }
}
return Trees;
```

Besides the problems mentioned above, this algorithm also has the problem that the sharing of subtrees is a non-trivial task. In fact, we assume that `ApplyTransformations` produces modified copies of T . To see how `ExhaustiveTransformation` works, consider again Figure 2.6. Assume that the top-left join tree is the initial join tree. Then, from this join tree `ApplyTransformations` produces all trees reachable by some edge. All of these are then added to `ToDo`. The next call to `ApplyTransformations` with any to the produced join trees will have the initial join tree contained in `Trees`. The complete set of visited join trees after this step is determined from the initial join tree by following at most two edges.

Let us reformulate the algorithm such that it uses a data structure similar to dynamic programming or memoization in order to avoid duplicate work. For any subset of relations, dynamic programming remembers the best join tree. This does not quite suffice for the transformation-based approach. Instead, we have to keep all join trees generated so far including those differing in the order of the arguments or a join operator. However, subtrees can be shared. This is done by keeping pointers into the data structure (see below). So, the difference between dynamic programming and the transformation-based approach becomes smaller. The main remaining difference is that dynamic programming only considers these join trees while with the transformation-based approach we have to keep the considered join trees since other join trees (more beneficial) might be generatable from them.

The data structure used for remembering trees is often called the MEMO structure. For every subset of relations to be joined (except the empty set), a *class* exists in the MEMO structure. Each class contains all the join trees that join exactly the relations describing the class. Here is an example for join trees containing three relations.

$\{R_1, R_2, R_3\}$	$\{R_1, R_2\} \bowtie R_3, R_3 \bowtie \{R_1, R_2\},$ $\{R_1, R_3\} \bowtie R_2, R_2 \bowtie \{R_1, R_3\},$ $\{R_2, R_3\} \bowtie R_1, R_1 \bowtie \{R_2, R_3\}$
$\{R_2, R_3\}$	$\{R_2\} \bowtie \{R_3\}, \{R_3\} \bowtie \{R_2\}$
$\{R_1, R_3\}$	$\{R_1\} \bowtie \{R_3\}, \{R_3\} \bowtie \{R_1\}$
$\{R_1, R_2\}$	$\{R_1\} \bowtie \{R_2\}, \{R_2\} \bowtie \{R_1\}$
$\{R_3\}$	R_3
$\{R_2\}$	R_2
$\{R_1\}$	R_1

Here, we used the set notation $\{\dots\}$ as an argument to a join to denote a reference to the class of join trees joining the relations contained in it.

We reformulate our transformation-based algorithm such that it fills in and uses the MEMO structure [581]. In a first step, the MEMO structure is initialized by creating an arbitrary join tree for the class $\{R_1, \dots, R_n\}$ and then going down this join tree and creating an entry for every join encountered. Then, we call `ExploreClass` on the root class comprising all relations to be joined. `ExploreClass` then applies `ApplyTransformations2` to every member of the class it is called upon. `ApplyTransformations2` then applies all rules to generate alternatives.

`ExhaustiveTransformation2(Query Graph G)`

Input: a query specification for relations $\{R_1, \dots, R_n\}$.

Output: an optimal join tree

```

initialize MEMO structure
ExploreClass( $\{R_1, \dots, R_n\}$ )
return best of class  $\{R_1, \dots, R_n\}$ 

```

`ExploreClass(C)`

Input: a class $C \subseteq \{R_1, \dots, R_n\}$

Output: none, but has side-effect on MEMO-structure

```

while (not all join trees in  $C$  have been explored) {
  choose an unexplored join tree  $T$  in  $C$ 
  ApplyTransformation2( $T$ )
  mark  $T$  as explored
}
return

```

`ApplyTransformations2(T)`

Input: a join tree of a class C

Output: none, but has side-effect on MEMO-structure

```

ExploreClass(left-child( $T$ ));
ExploreClass(right-child( $T$ ));
foreach transformation  $\mathcal{T}$  and class member of child classes {

```

```

foreach  $T'$  resulting from applying  $T$  to  $T$  {
  if  $T'$  not in MEMO structure {
    add  $T'$  to class  $C$  of MEMO structure
  }
}
return

```

`ApplyTransformations2` uses a set of transformations to be applied. We discuss now the effect of different transformation sets on the complexity of the algorithm. Applying `ExhaustiveTransformation2` with a rule set consisting of Commutativity and Left and Right Associativity generates $4^n - 3^{n+1} + 2^{n+2} - n - 2$ duplicates for n relations. Contrast this with the number of join trees contained in a completely filled MEMO structure³: $3^n - 2^{n+1} + n + 1$. This clearly shows the problem.

The problem of generating the same join tree several times was considered by Pellenkoft, Galindo-Legaria, and Kersten [581, 582, 583]. The solution lies in parameterizing `ExhaustiveTransformation2` by an appropriate set of transformations. The basic idea is to remember for every join operator which rules are applicable to it. For example, after applying commutativity to a join operator, we disable commutativity for it.

For acyclic queries, the following rule set guarantees that all bushy join trees are generated, but no duplicates [583]. Thereby, cross products are not considered. That is, a rule is only applicable if it does not result in a cross product. This restricts the applicability of the above algorithm to connected queries. We use C_i to denote some class of the MEMO structure. We call the following rule set RS-1:

RS-1

T_1 : Commutativity $C_1 \bowtie_0 C_2 \rightsquigarrow C_2 \bowtie_1 C_1$

Disable all transformations T_1 , T_2 , and T_3 for \bowtie_1 .

T_2 : Right Associativity $(C_1 \bowtie_0 C_2) \bowtie_1 C_3 \rightsquigarrow C_1 \bowtie_2 (C_2 \bowtie_3 C_3)$

Disable transformations T_2 and T_3 for \bowtie_2 and enable all rules for \bowtie_3 .

T_3 : Left associativity $C_1 \bowtie_0 (C_2 \bowtie_1 C_3) \rightsquigarrow (C_1 \bowtie_2 C_2) \bowtie_3 C_3$

Disable transformations T_2 and T_3 for \bowtie_3 and enable all rules for \bowtie_2 .

In order to be able to follow these rules, the procedure `ApplyTransformations2` has to be enhanced such that it is able to keep track of the application history of the rules for every join operator. The additional memory requirement is neglectable, since a single bit for each rules suffices.

As an example, let us consider the chain query $R_1 - R_2 - R_3 - R_4$. Figure 3.14 shows the MEMO structure. The first column gives the sets of the relations identifying each class. We leave out the single relation classes assuming that $\{R_i\}$ has R_i as its only join tree which is marked as explored.

³The difference to the according number for dynamic programming is due to the fact that we have to keep alternatives generated by commutativity and that join trees for single relations are counted.

Class	Initialization	Transformation	Step
$\{R_1, R_2, R_3, R_4\}$	$\{R_1, R_2\} \bowtie_{111} \{R_3, R_4\}$	$\{R_3, R_4\} \bowtie_{000} \{R_1, R_2\}$	3
		$R_1 \bowtie_{100} \{R_2, R_3, R_4\}$	4
		$\{R_1, R_2, R_3\} \bowtie_{100} R_4$	5
		$\{R_2, R_3, R_4\} \bowtie_{000} R_1$	8
		$R_4 \bowtie_{000} \{R_1, R_2, R_3\}$	10
$\{R_2, R_3, R_4\}$		$R_2 \bowtie_{111} \{R_3, R_4\}$	4
		$\{R_3, R_4\} \bowtie_{000} R_2$	6
		$\{R_2, R_3\} \bowtie_{100} R_4$	6
		$R_4 \bowtie_{000} \{R_2, R_3\}$	7
$\{R_1, R_3, R_4\}$			
$\{R_1, R_2, R_4\}$			
$\{R_1, R_2, R_3\}$			
		$\{R_1, R_2\} \bowtie_{111} R_3$	5
		$R_3 \bowtie_{000} \{R_1, R_2\}$	9
		$R_1 \bowtie_{100} \{R_2, R_3\}$	9
		$\{R_2, R_3\} \bowtie_{000} R_1$	9
$\{R_3, R_4\}$	$R_3 \bowtie_{111} R_4$	$R_4 \bowtie_{000} R_3$	2
$\{R_2, R_4\}$			
$\{R_2, R_3\}$			
$\{R_1, R_4\}$			
$\{R_1, R_3\}$			
$\{R_1, R_2\}$	$R_1 \bowtie_{111} R_2$	$R_2 \bowtie_{000} R_1$	1

Figure 3.14: Example of rule transformations (RS-1)

The second column shows the initialization with an arbitrarily chosen join tree. The third column is the one filled by the `Apply Transformation2` procedure. We apply the rule set RS-1, which consists of three transformations. Each join is annotated with three bits, where the i -th bit indicates whether T_i is applicable (1) or not (0). After initializing the MEMO structure, `ExhaustiveTransformation2` calls `ExploreClass` for $\{R_1, R_2, R_3, R_4\}$. The only (unexplored) join tree is $\{R_1, R_2\} \bowtie_{111} \{R_3, R_4\}$, which will become the argument of `ApplyTransformations2`. Next, `ExploreClass` is called on $\{R_1, R_2\}$ and $\{R_3, R_4\}$. In both cases, T_1 is the only applicable rule, and the result is shown in the third column under steps 1 and 2. Now we have to apply all transformations on $\{R_1, R_2\} \bowtie_{111} \{R_3, R_4\}$. Commutativity T_1 gives us $\{R_3, R_4\} \bowtie_{000} \{R_1, R_2\}$ (Step 3). For right associativity, we have two elements in class $\{R_1, R_2\}$. Substituting them and applying T_2 gives

1. $(R_1 \bowtie R_2) \bowtie \{R_3, R_4\} \rightsquigarrow R_1 \bowtie_{100} (R_2 \bowtie_{111} \{R_3, R_4\})$
2. $(R_2 \bowtie R_1) \bowtie \{R_3, R_4\} \rightsquigarrow R_2 \bowtie_{111} (R_1 \times \{R_3, R_4\})$

The latter contains a cross product. This leaves us with the former as the result of Step 4. The right argument of the top most join is $R_2 \bowtie_{111} \{R_3, R_4\}$. Since

we do not find it in class $\{R_2, R_3, R_4\}$, we add it (4).

T_3 is next.

1. $\{R_1, R_2\} \bowtie (R_3 \bowtie R_4) \rightsquigarrow (\{R_1, R_2\} \bowtie_{111} R_3) \bowtie_{100} R_4$
2. $\{R_1, R_2\} \bowtie (R_4 \bowtie R_3) \rightsquigarrow (\{R_1, R_2\} \times R_4) \bowtie_{100} R_3$

The latter contains a cross product. This leaves us with the former as the result of Step 5. We also add $\{R_1, R_2\} \bowtie_{111} R_3$. Now that $\{R_1, R_2\} \bowtie_{111} \{R_3, R_4\}$ is completely explored, we turn to $\{R_3, R_4\} \bowtie_{000} \{R_1, R_2\}$, but all transformations are disabled here.

$R_1 \bowtie_{100} \{R_2, R_3, R_4\}$ is next. First, $\{R_2, R_3, R_4\}$ has to be explored. The only entry is $R_2 \bowtie_{111} \{R_3, R_4\}$. Remember that $\{R_3, R_4\}$ is already explored. T_2 is not applicable. The other two transformations give us

$$T_1 \{R_3, R_4\} \bowtie_{000} R_2$$

$$T_3 (R_2 \bowtie_{000} R_3) \bowtie_{100} R_4 \text{ and } (R_2 \times R_4) \bowtie_{100} R_3$$

Those join trees not exhibiting a cross product are added to the MEMO structure under 6. Applying commutativity to $\{R_2, R_4\} \bowtie_{100} R_3$ gives 7. Commutativity is the only rule enabled for $R_1 \bowtie_{100} \{R_2, R_3, R_4\}$. Its application results in 8.

$\{R_1, R_2, R_3\} \bowtie_{100} R_4$ is next. It is simple to explore the class $\{R_1, R_2, R_3\}$ with its only entry $\{R_1, R_2\} \bowtie_{111} R_3$:

$$T_1 R_3 \bowtie_{000} \{R_1, R_2\}$$

$$T_2 R_1 \bowtie_{100} (R_2 \bowtie_{111} R_3) \text{ and } R_2 \bowtie_{100} (R_1 \times R_3)$$

Commutativity can still be applied to $R_1 \bowtie_{100} (R_2 \bowtie_{111} R_3)$. All the new entries are numbered 9. Commutativity is the only rule enabled for $\{R_1, R_2, R_3\} \bowtie_{100} R_4$. Its application results in 10.

□

The next two sets of transformations were originally intended for generating all bushy/left-deep trees for a clique query [582]. They can, however, also be used to generate all bushy trees when cross products are considered. The rule set RS-2 for bushy trees is

$$T_1: \textbf{Commutativity} \ C_1 \bowtie_0 C_2 \rightsquigarrow C_2 \bowtie_1 C_1$$

Disable all transformations $T_1, T_2, T_3,$ and T_4 for \bowtie_1 .

$$T_2: \textbf{Right Associativity} \ (C_1 \bowtie_0 C_2) \bowtie_1 C_3 \rightsquigarrow C_1 \bowtie_2 (C_2 \bowtie_3 C_3)$$

Disable transformations $T_2, T_3,$ and T_4 for \bowtie_2 .

$$T_3: \textbf{Left Associativity} \ C_1 \bowtie_0 (C_2 \bowtie_1 C_3) \rightsquigarrow (C_1 \bowtie_2 C_2) \bowtie_3 C_3$$

Disable transformations T_2, T_3 and T_4 for \bowtie_3 .

$$T_4: \textbf{Exchange} \ (C_1 \bowtie_0 C_2) \bowtie_1 (C_3 \bowtie_2 C_4) \rightsquigarrow (C_1 \bowtie_3 C_3) \bowtie_4 (C_2 \bowtie_5 C_4)$$

Disable all transformations $T_1, T_2, T_3,$ and T_4 for \bowtie_4 .

If we initialize the MEMO structure with left-deep trees, we can strip down the above rule set to Commutativity and Left Associativity. The reason is an observation made by Shapiro et al.: from a left-deep join tree we can generate all bushy trees with only these two rules [688].

If we want to consider only left-deep trees, the following rule set RS-3 is appropriate:

T_1 **Commutativity** $R_1 \bowtie_0 R_2 \rightsquigarrow R_2 \bowtie_1 R_1$

Here, the R_i are restricted to classes with exactly one relation. T_1 is disabled for \bowtie_1 .

T_2 **Right Join Exchange** $(C_1 \bowtie_0 C_2) \bowtie_1 C_3 \rightsquigarrow (C_1 \bowtie_2 C_3) \bowtie_3 C_2$

Disable T_2 for \bowtie_3 .

3.3 Probabilistic Algorithms

3.3.1 Generating Random Left-Deep Join Trees with Cross Products

The basic idea of the algorithms in this section and the following sections is to generate a set of randomly chosen join trees, evaluate their costs, and return the best one. The problem with this approach lies in the random generation of join trees: every join tree has to be generated with equal probability. Although there are some advocates of the pure random approach [258, 259, 261, 257], typically a random join tree or a set of random join trees is used in subsequent algorithms like iterative improvement and simulated annealing.

Obviously, if we do not consider cross products the problem is really hard, since the query graph plays an important role. So let us start with the simplest case where random join trees are generated that might contain cross products even for connected query graphs. Then, any join tree is a valid join tree.

The general idea behind all algorithms is the following. Assume that the number of join trees in the considered search space is known to be N . Then, instead of generating a random join tree directly, a bijective mapping from the interval of non-negative integers $[0, N[$ to a join tree in the search space is established. Then, a random join tree can be generated by (1) generating a random number in $[0, N[$ and (2) mapping the number to the join tree. The problem of bijectively mapping an interval of non-negative integers to elements of a set is usually called *unranking*. The opposite mapping is called *ranking*. Obviously, the crux in our case is the efficiency of the unranking problem.

We start with generating random left-deep join trees for n relations. This problem is identical to generating random permutations. That is, we look for a fast unranking algorithm that maps the non-negative integers in $[0, n[$ to permutations. Let us consider permutations of the numbers $\{0, \dots, n-1\}$. A mapping between these numbers and relations is established easily, e.g. via an array. The traditional approach to ranking/unranking of permutations is to first define an ordering on the permutations and then find a ranking and unranking algorithm relative to that ordering. For the lexicographic order, algorithms require $O(n^2)$ time [479, 617]. More sophisticated algorithms separate

the ranking/unranking algorithms into two phases. For ranking, first the *inversion vector* of the permutation is established. Then, ranking takes place for the inversion vector. Unranking works in the opposite direction. The *inversion vector* of a permutation $\pi = \pi_0, \dots, \pi_{n-1}$ is defined to be the sequence $v = v_0, \dots, v_{n-1}$, where v_i is equal to the number of entries π_j with $\pi_j > \pi_i$ and $j < i$. Inversion vectors uniquely determine a permutation [759]. However, naive algorithms of this approach again require $O(n^2)$ time. Better algorithms require $O(n \log n)$. Using an elaborated data structure, Dietz' algorithm requires $O((n \log n)/(\log \log n))$ [208]. Other orders like the Steinhaus-Johnson-Trotter order have been exploited for ranking/unranking but do not yield any run-time advantage over the above mentioned algorithms (see [442, 617]).

Since it is not important for our problem that any order constraints are satisfied for the ranking/unranking functions, we use the fastest possible algorithm established by Myrvold and Ruskey [541]. It runs in $O(n)$ which is also easily seen to be a lower bound.

The algorithm is based on the standard algorithm to generate random permutations [192, 214, 535]. An array π is initialized such that $\pi[i] = i$ for $0 \leq i \leq n - 1$. Then, the loop

```
for (k = n - 1; k >= 0; -- k) swap( $\pi[k]$ ,  $\pi[\text{random}(k)]$ );
```

is executed where `swap` exchanges two elements and `random(k)` generates a random number in $[0, k]$. This algorithm randomly picks any of the possible permutations. Assume the random elements produced by the algorithm are r_{n-1}, \dots, r_0 where $0 \leq r_i \leq i$. Obviously, there are exactly $n(n-1)(n-2) \dots 1 = n!$ such sequences and there is a one-to-one correspondence between these sequences and the set of all permutations. We can thus unrank $r \in [0, n!$ by turning it into a unique sequence of values r_{n-1}, \dots, r_0 . Note that after executing the swap with r_{n-1} , every value in $[0, n[$ is possible at position $\pi[n-1]$. Further, $\pi[n-1]$ is never touched again. Hence, we can unrank r as follows. We first set $r_{n-1} = r \bmod n$ and perform the swap. Then, we define $r' = \lfloor r/n \rfloor$ and iteratively unrank r' to construct a permutation of $n-1$ elements. The following algorithm realizes this idea.

```
Unrank(n, r) {
  Input: the number n of elements to be permuted
         and the rank r of the permutation to be constructed
  Output: a permutation  $\pi$ 
  for (i = 0; i < n; ++ i)  $\pi[i] = i$ ;
  Unrank-Sub(n, r,  $\pi$ );
  return  $\pi$ ;
}
```

```
Unrank-Sub(n, r,  $\pi$ ) {
  for (i = n; i > 0; -- i) {
```

```

    swap( $\pi[i - 1]$ ,  $\pi[r \bmod i]$ );
     $r = \lfloor r/i \rfloor$ ;
  }
}
```

3.3.2 Generating Random Join Trees with Cross Products

Next, we want to randomly construct bushy plans possibly containing cross products. This is done in several steps:

1. Generate a random number b in $[0, C(n - 1)[$.
2. Unrank b to obtain a bushy tree with $n - 1$ inner nodes.
3. Generate a random number p in $[0, n![$.
4. Unrank p to obtain a permutation.
5. Attach the relations in order p from left to right as leaf nodes to the binary tree obtained in Step 2.

The only step that we still have to discuss is Step 2. It is a little involved and we can only try to bring across the general idea. For details, the reader is referred to the literature [479, 480, 481].

Consider Figure 3.15. It contains all 14 possible trees with four inner nodes. The trees are ordered according to the rank we will consider. The bottom-most number below any tree is its rank in $[0, 14[$. While unranking, we do not generate the trees directly, but an encoding of the tree instead. This encoding works as follows. Any binary tree corresponds to a word in a Dyck language with one pair of parenthesis. The alphabet hence consists of $\Sigma = \{ '(', ') '\}$. For join trees with n inner nodes, we use Dyck words of length $2n$ whose parenthesization is correct. That is, for every $'('$, we have a subsequent $)'$. From a given join tree, we obtain the Dyck word by a preorder traversal. Whenever we encounter an inner node, we encode this with a $'('$. All but the last leaf nodes are encoded by a $)'$. Appending all these $2n$ encodings gives us a Dyck word of length $2n$. Figure 3.15 shows directly below each tree its corresponding Dyck word. In the line below, we simply changed the representation by substituting every $'('$ by a $'1'$ and every $)'$ by a $'0'$. The encoding that will be generated by the unranking algorithm is shown in the third line below each tree: we remember the places (index in the bit-string) where we find a $'1'$.

In order to do the unranking, we need to do some counting. Therefor, we map Dyck words to paths in a triangular grid. For $n = 4$ this grid is shown in Figure 3.16. We always start at $(0, 0)$ which means that we have not opened a parenthesis. When we are at (i, j) , opening a parenthesis corresponds to going to $(i + 1, j + 1)$ and closing a parenthesis to going to $(i + 1, j - 1)$. We have thus established a bijective mapping between Dyck words and paths in the grid. Thus counting Dyck words corresponds to counting paths.




Figure 3.15: Encoding Trees

The number of different paths from $(0, 0)$ to (i, j) can be computed by

$$p(i, j) = \frac{j+1}{i+1} \binom{i+1}{\frac{1}{2}(i+j)+1}$$

These numbers are called the *Ballot numbers* [108]. The number of paths from (i, j) to $(2n, 0)$ can thus be computed as (see [480, 481]):

$$q(i, j) = p(2n - i, j)$$




Figure 3.16: Paths

Note the special case $q(0,0) = p(2n,0) = C(n)$. In Figure 3.16, we annotated nodes (i, j) by $p(i, j)$. These numbers can be used to assign (sub-) intervals to paths (Dyck words, trees). For example, if we are at $(4, 4)$, there exists only a single path to $(2n, 0)$. Hence, the path that travels the edge $(4, 4) \rightarrow (5, 3)$ has rank 0. From $(3, 3)$ there are four paths to $(2n, 0)$, one of which we already considered. This leaves us with three paths that travel the edge $(3, 3) \rightarrow (4, 2)$. The paths in this part as assigned ranks in the interval $[1, 4[$. Figure 3.16 shows the intervals near the edges. For unranking, we can now proceed as follows. Assume we have a rank r . We consider opening a parenthesis (go from (i, j) to $(i + 1, j + 1)$) as long as the number of paths from that point does no longer exceed our rank r . If it does, we close a parenthesis instead (go from (i, j) to $(i - 1, j + 1)$). Assume, that we went upwards to (i, j) and then had to go down to $(i - 1, j + 1)$. We subtract the number of paths from $(i + 1, j + 1)$ from our rank r and proceed iteratively from $(i - 1, j + 1)$ by going up as long as possible and going down again. Remembering the number of parenthesis opened and closed along our way results in the required encoding. The following algorithm finalizes these ideas.

```
UnrankTree(n, r)
```

Input: a number of inner nodes n and a rank $r \in [0, C(n - 1)]$

Output: encoding of the inner leaves of a tree

```
lNoParOpen = 0;
```

```
lNoParClose = 0;
```

```
i = 1; // current encoding
```

```
j = 0; // current position in encoding array
```

```
while (j < n) {
```

```
    k = q(lNoParOpen + lNoParClose + 1, lNoParOpen - lNoParClose + 1);
```

```
    if (k ≤ r) {
```

```
        r -= k;
```

```

    ++lNoParClose;
  } else {
    aTreeEncoding[j++] = i;
    ++lNoParOpen;
  }
  ++i;
}

```

Given an array with the encoding of a tree, it is easy to construct the tree from it. The following procedure does that.

```

TreeEncoding2Tree(n, aEncoding) {
Input: the number of internal nodes of the tree n
Output: root node of the result tree
  root = new Node; /* root of the result tree */
  curr = root; /* curr: current internal node whose subtrees are to be created */
  i = 1; /* pointer to entry in encoding */
  child = 0; /* 0 = left , 1 = right: next child whose subtree is to be created */
  while (i < n) {
    lDiff = aEncoding[i] - aEncoding[i - 1];
    for (k = 1; k < lDiff; ++k) {
      if (child == 0) {
        curr->addLeftLeaf();
        child = 1;
      } else {
        curr->addRightLeaf();
        while (curr->right() != 0) {
          curr = curr->parent();
        }
        child = 1;
      }
    }
  }
  if (child == 0) {
    curr->left(new Node(curr)); // curr becomes parent of new node
    curr = curr->left();
    ++i;
    child = 0;
  } else {
    curr->right(new Node(curr));
    curr = curr->right();
    ++i;
    child = 0;
  }
}
while (curr != 0) {

```

```

    curr->addLeftLeaf(); // addLeftLeaf adds leaf if no left-child exists
    curr->addRightLeaf(); // analogous
    curr = curr->parent();
}
return root;
}

```

3.3.3 Generating Random Join Trees without Cross Products

A general solution for randomly generating join trees without cross products is not known. However, if we restrict ourselves to acyclic queries, we can apply an algorithm developed by Galindo-Legaria, Pellenkofft, and Kersten [259, 258, 261]. For this algorithm to work, we have to assume that the query graph is connected and acyclic.

For the rest of this section, we assume that $G = (V, E)$ is the query graph and $|V| = n$. That is, n relations are to be joined. No join tree contains a cross product. With every node in a join tree, we associate a *level*. The root has level 0. Its children have level 1, and so on. We further use lower-case letters for relations.

For a given query graph G , we denote by \mathcal{T}_G the set of join trees for G . Let $\mathcal{T}_G^{v(k)} \subseteq \mathcal{T}_G$ be the subset of join trees where the leaf node (i.e. relation) v occurs at level k . Some trivial observations follow. If the query graph consists of a single node ($n = 1$), then $|\mathcal{T}_G| = |\mathcal{T}_G^{v(0)}| = 1$. If $n > 1$, the top node in the join tree is a join and not a relation. Hence, $|\mathcal{T}_G^{v(0)}| = 0$. Obviously, the maximum level that can occur in any join tree is $n - 1$. Hence, $|\mathcal{T}_G^{v(k)}| = 0$ if $k \geq n$. Since the level at which a leaf node v occurs in some join tree is unique, we have $\mathcal{T}_G = \cup_{k=0}^n \mathcal{T}_G^{v(k)}$ and $\mathcal{T}_G^{v(i)} \cap \mathcal{T}_G^{v(j)} = \emptyset$ for $i \neq j$. This gives us $|\mathcal{T}_G| = \sum_{k=0}^n |\mathcal{T}_G^{v(k)}|$.

The algorithm generates an unordered tree with n leaf nodes. If we wish to have a random ordered tree, we have to pick one of the 2^{n-1} possibilities to order the $(n - 1)$ joins within the tree. We proceed as follows. We start with some notation for lists, discuss how two lists can be merged, describe how a specific merge can be specified, and count the number of possible merges. This is important, since join trees will be described as lists of trees. Given a leaf node v , we simply traverse the path from the root to v . Thereby, subtrees that branch off can be collected into a list of trees. After these remarks, we start developing the algorithm in several steps. First, we consider two operations with which we can construct new join trees: *leaf-insertion* introduces a new leaf node into a given tree and *tree-merging* merges two join trees. Since we do not want to generate cross products in this section, we have to apply these operations carefully. Therefor, we need a description of how to generate *all* valid join trees for a given query graph. The central data structure for this purpose is the *standard decomposition graph* (SDG). Hence, in the second step, we define SDGs and introduce an algorithm that derives an SDG from a given query graph. In the third step, we start counting. The fourth and final step consists of the unranking algorithm. We do not discuss the ranking algorithm.




Figure 3.17: Tree-merge

It can be found in [261].

We use the Prolog notation $|$ to separate the first element of a list from its tail. For example, the list $\langle a|t \rangle$ has a as its first element and a tail t . Assume that P is a property of elements. A list l' is the *projection* of a list L on P , if L' contains all elements of L satisfying the property P . Thereby, the order is retained. A list L is a *merge* of two disjoint lists L_1 and L_2 if L contains all elements from L_1 and L_2 and both are projections of L .

A merge of a list L_1 with a list L_2 whose respective lengths are l_1 and l_2 can be described by an array $\alpha = [\alpha_0, \dots, \alpha_{l_2}]$ of non-negative integers whose sum is equal to l_1 . The non-negative integer α_{i-1} gives the number of elements of L_1 which precede the i -th element of L_2 in the merged list. We obtain the merged list L by first taking α_0 elements from L_1 . Then, an element from L_2 follows. Then α_1 elements from L_1 and the next element of L_2 follow and so on. Finally follow the last α_{l_2} elements of L_1 . Figure 3.17 illustrates possible merges.

Compare list merges to the problem of non-negative (weak) integer composition [?]. There, we ask for the number of compositions of a non-negative integer n into k non-negative integers α_i with $\sum_{i=1}^k \alpha_k = n$. The answer is $\binom{n+k-1}{k-1}$ [718]. Since we have to decompose l_1 into $l_2 + 1$ non-negative integers, the number of possible merges is $M(l_1, l_2) = \binom{l_1+l_2}{l_2}$. The observation $M(l_1, l_2) = M(l_1 - 1, l_2) + M(l_1, l_2 - 1)$ allows us to construct an array of size $n * n$ in $O(n^2)$ that materializes the values for M . This array will allow us to rank list merges in $O(l_1 + l_2)$.

The idea for establishing a bijection between $[1, M(l_1, l_2)]$ and the possible α s is a general one and used for all subsequent algorithms of this section. Assume that we want to rank the elements of some set S and $S = \cup_{i=0}^n S_i$ is partitioned into disjoint S_i . If we want to rank $x \in S_k$, we first find the *local*

rank of $x \in S_k$. The rank of x is then defined as

$$\sum_{i=0}^{k-1} |S_i| + \text{local-rank}(x, S_k)$$

To unrank some number $r \in [1, N]$, we first find k such that

$$k = \min_j (r \leq \sum_{i=0}^j |S_i|)$$

Then, we proceed by unranking with the new local rank

$$r' = r - \sum_{i=0}^{k-1} |S_i|$$

within S_k .

Accordingly, we partition the set of all possible merges into subsets. Each subset is determined by α_0 . For example, the set of possible merges of two lists L_1 and L_2 with length $l_1 = l_2 = 4$ is partitioned into subsets with $\alpha_0 = j$ for $0 \leq j \leq 4$. In each partition, we have $M(j, l_2 - 1)$ elements. To unrank a number $r \in [1, M(l_1, l_2)]$, we first determine the partition by computing $k = \min_j r \leq \sum_{i=0}^j M(j, l_2 - 1)$. Then, $\alpha_0 = l_1 - k$. With the new rank $r' = r - \sum_{i=0}^k M(j, l_2 - 1)$, we start iterating all over. The following table gives the numbers for our example and can be used to understand the unranking algorithm. The algorithm itself can be found in Figure 3.18.

k	α_0	$(k, l_2 - 1)$	$M(k, l_2 - 1)$	rank intervals
0	4	(0, 3)	1	[1, 1]
1	3	(1, 3)	4	[2, 5]
2	2	(2, 3)	10	[6, 15]
3	1	(3, 3)	20	[16, 35]
4	0	(4, 3)	35	[36, 70]

We now turn to the *anchored list representation* of join trees.

Definition 3.3.1 *Let T be a join tree and v be a leaf of T . The anchored list representation L of T is constructed as follows:*

- *If T consists of the single leaf node v , then $L = \langle \rangle$.*
- *If $T = (T_1 \bowtie T_2)$ and without loss of generality v occurs in T_2 , then $L = \langle T_1 | L_2 \rangle$, where L_2 is the anchored list representation of T_2 .*

We then write $T = (L, v)$.

Observe that if $T = (L, v) \in \mathcal{T}_G$, then $T \in \mathcal{T}_G^{v(k)} \prec \succ |L| = k$.

The operation *leaf-insertion* is illustrated in Figure 3.19. A new leaf v is inserted into the tree at level k . Formally, it is defined as follows.

```

UnrankDecomposition( $r, l_1, l_2$ )
Input: a rank  $r$ , two list sizes  $l_1$  and  $l_2$ 
Output: a merge specification  $\alpha$ .
for ( $i = 0; i \leq l_2; ++i$ ) {
     $\alpha[i] = 0;$ 
}
 $i = k = 0;$ 
while ( $l_1 > 0 \ \&\& \ l_2 > 0$ ) {
     $m = M(k, l_2 - 1);$ 
    if ( $r \leq m$ ) {
         $\alpha[i++] = l_1 - k;$ 
         $l_1 = k;$ 
         $k = 0;$ 
         $--l_2;$ 
    } else {
         $r -= m;$ 
         $++k;$ 
    }
}
 $\alpha[i] = l_1;$ 
return  $\alpha;$ 

```

Figure 3.18: Algorithm UnrankDecomposition




Figure 3.19: Leaf-insertion

Definition 3.3.2 Let $G = (V, E)$ be a query graph, T a join tree of G . $v \in V$ be such that $G' = G|_{V \setminus \{v\}}$ is connected, $(v, w) \in E$, $1 \leq k < n$, and

$$T = (\langle T_1, \dots, T_{k-1}, v, T_{k+1}, \dots, T_n \rangle, w) \quad (3.10)$$

$$T' = (\langle T_1, \dots, T_{k-1}, T_{k+1}, \dots, T_n \rangle, w). \quad (3.11)$$




Figure 3.20: A query graph, its tree, and its standard decomposition graph

Then we call (T', k) an insertion pair on v and say that T is decomposed into (or constructed from) the pair (T', k) on v .

Observe that leaf-insertion defines a bijective mapping between $\mathcal{T}_G^{v(k)}$ and insertion pairs (T', k) on v , where T' is an element of the disjoint union $\cup_{i=k-1}^{n-2} \mathcal{T}_{G'}^{w(i)}$.

The operation *tree-merging* is illustrated in Figure 3.17. Two trees $R = (L_R, w)$ and $S = (L_S, w)$ on a common leaf w are merged by merging their anchored list representations.

Definition 3.3.3 Let $G = (V, E)$ be a query graph, $w \in V$, $T = (L, w)$ a join tree of G , $V_1, V_2 \subseteq V$ such that $G_1 = G|_{V_1}$ and $G_2 = G|_{V_2}$ are connected, $V_1 \cup V_2 = V$, and $V_1 \cap V_2 = \{w\}$. For $i = 1, 2$:

- Define the property P_i to be “every leaf of the subtree is in V_i ”,
- Let L_i be the projection of L on P_i .
- $T_i = (L_i, w)$.

Let α be the integer composition such that L is the result of merging L_1 and L_2 on α . Then we call (T_1, T_2, α) a merge triplet. We say that T is decomposed into (constructed from) (T_1, T_2, α) on V_1 and V_2 .

Observe that the *tree-merging* operation defines a bijective mapping between $\mathcal{T}_G^{w(k)}$ and merge triplets (T_1, T_2, α) , where $T_1 \in \mathcal{T}_{G_1}^{w(i)}$, $T_2 \in \mathcal{T}_{G_2}^{w(k-i)}$, and α specifies a merge of two lists of sizes i and $k - i$. Further, the number of these merges (i.e. the number of possibilities for α) is $\binom{i+(k-i)}{k-i} = \binom{k}{i}$.

A *standard decomposition graph* of a query graph describes the possible constructions of join trees. It is not unique (for $n > 1$) but anyone can be used to construct all possible unordered join trees. For each of our two operations it has one kind of inner nodes. A unary node labeled $+_v$ stands for leaf-insertion of v . A binary node labeled $*_w$ stands for tree-merging its subtrees whose only common leaf is w .

The standard decomposition graph of a query graph $G = (V, E)$ is constructed in three steps:

1. pick an arbitrary node $r \in V$ as its root node;
2. transform G into a tree G' by directing all edges away from r ;
3. call $\text{QG2SDG}(G', r)$

with

```

QG2SDG( $G', r$ )
Input: a query tree  $G' = (V, E)$  and its root  $r$ 
Output: a standard query decomposition tree of  $G'$ 
Let  $\{w_1, \dots, w_n\}$  be the children of  $v$ ;
switch ( $n$ ) {
  case 0: label  $v$  with " $v$ ";
  case 1:
    label  $v$  as " $+v$ ";
    QG2SDG( $G', w_1$ );
  otherwise:
    label  $v$  as " $*v$ ";
    create new nodes  $l, r$  with label  $+v$ ;
     $E \setminus = \{(v, w_i) | 1 \leq i \leq n\}$ ;
     $E \cup = \{(v, l), (v, r), (l, w_1)\} \cup \{(r, w_i) | 2 \leq i \leq n\}$ ;
    QG2SDG( $G', l$ );
    QG2SDG( $G', r$ );
}
return  $G'$ ;

```

Note that QG2SDG transforms the original graph G' into its SDG by side-effects. Thereby, the n -ary tree is transformed into a binary tree similar to the procedure described by Knuth [433, Chap 2.3.2]. Figure 3.20 shows a query graph G , its tree G' rooted at e , and its standard decomposition tree.

For an efficient access to the number of join trees in some partition $\mathcal{T}_G^{v(k)}$ in the unranking algorithm, we materialize these numbers. This is done in the `count` array. The semantics of a `count` array $[c_0, c_1, \dots, c_n]$ of a node u with label \circ_v ($\circ \in \{+, *\}$) of the SDG is that u can construct c_i different trees in which leaf v is at level i . Then, the total number of trees for a query can be computed by summing up all the c_i in the `count` array of the root node of the decomposition tree.

To compute the `count` and an additional `summand` adornment of a node labeled $+v$, we use the following lemma.

Lemma 3.3.4 *Let $G = (V, E)$ be a query graph with n nodes, $v \in V$ such that $G' = G|_{V \setminus v}$ is connected, $(v, w) \in E$, and $1 \leq k < n$. Then*

$$|\mathcal{T}_G^{v(k)}| = \sum_{i \geq k-1} |\mathcal{T}_{G'}^{w(i)}|$$

This lemma follows from the observation made after the definition of the leaf-insertion operation.

The sets $\mathcal{T}_{G'}^{w(i)}$ used in the summands of Lemma 3.3.4 directly correspond to subsets $\mathcal{T}_G^{v(k),i}$ ($k-1 \leq i \leq n-2$) defined such that $T \in \mathcal{T}_G^{v(k),i}$ if

1. $T \in \mathcal{T}_G^{v(k)}$,
2. the insertion pair on v of T is (T', k) , and
3. $T' \in \mathcal{T}_{G'}^{w(i)}$.

Further, $|\mathcal{T}_G^{v(k),i}| = |\mathcal{T}_{G'}^{w(i)}|$. For efficiency, we materialize the summands in an array of arrays **summands**.

To compute the **count** and **summand** adornment of a node labeled $*_v$, we use the following lemma.

Lemma 3.3.5 *Let $G = (V, E)$ be a query graph, $w \in V$, $T = (L, w)$ a join tree of G , $V_1, V_2 \subseteq V$ such that $G_1 = G|_{V_1}$ and $G_2 = G|_{V_2}$ are connected, $V_1 \cup V_2 = V$, and $V_1 \cap V_2 = \{v\}$. Then*

$$|\mathcal{T}_G^{v(k)}| = \sum_i \binom{k}{i} |\mathcal{T}_{G_1}^{v(i)}| |\mathcal{T}_{G_2}^{v(k-i)}|$$

This lemma follows from the observation made after the definition of the tree-merge operation.

The sets $\mathcal{T}_{G'}^{w(i)}$ used in the summands of Lemma 3.3.5 directly correspond to subsets $\mathcal{T}_G^{v(k),i}$ ($0 \leq i \leq k$) defined such that $T \in \mathcal{T}_G^{v(k),i}$ if

1. $T \in \mathcal{T}_G^{v(k)}$,
2. the merge triplet on V_1 and V_2 of T is (T_1, T_2, α) , and
3. $T_1 \in \mathcal{T}_{G_1}^{v(i)}$.

Further, $|\mathcal{T}_G^{v(k),i}| = \binom{k}{i} |\mathcal{T}_{G_1}^{v(i)}| |\mathcal{T}_{G_2}^{v(k-i)}|$.

Before we come to the algorithm for computing the adornments **count** and **summands**, let us make one observation that follows directly from the above two lemmata. Assume a node v whose **count** array is $[c_1, \dots, c_m]$ and whose **summands** is $s = [s^0, \dots, s^n]$ with $s_i = [s_0^i, \dots, s_m^i]$, then $c_i = \sum_{j=0}^m s_j^i$ holds. Figure 3.21 contains the algorithm to adorn SDG's nodes with **count** and **summands**. It has worst-case complexity $O(n^3)$. Figure 3.20 shows the **count** adornment for the SDG. Looking at the **count** array of the root node, we see that the total number of join trees for our example query graph is 18.

The algorithm `UnrankLocalTreeNoCross` called by `UnrankTreeNoCross` adorns the standard decomposition graph with **insert-at** and **merge-using** annotations. These can then be used to extract the join tree.

```

Adorn( $v$ )
Input: a node  $v$  of the SDG
Output:  $v$  and nodes below are adorned by count and summands
Let  $\{w_1, \dots, w_n\}$  be the children of  $v$ ;
switch ( $n$ ) {
  case 0: count( $v$ ) := [1]; // no summands for  $v$ 
  case 1:
    Adorn( $w_1$ );
    assume count( $w_1$ ) = [ $c_0^1, \dots, c_{m_1}^1$ ];
    count( $v$ ) = [ $0, c_1, \dots, c_{m_1+1}$ ] where  $c_k = \sum_{i=k-1}^{m_1} c_i^1$ ;
    summands( $v$ ) = [ $s^0, \dots, s^{m_1+1}$ ] where  $s^k = [s_0^k, \dots, s_{m_1+1}^k]$  and
     $s_i^k = \begin{cases} c_i^1 & \text{if } 0 < k \text{ and } k-1 \leq i \\ 0 & \text{else} \end{cases}$ 
  case 2:
    Adorn( $w_1$ );
    Adorn( $w_2$ );
    assume count( $w_1$ ) = [ $c_0^1, \dots, c_{m_1}^1$ ];
    assume count( $w_2$ ) = [ $c_0^2, \dots, c_{m_2}^2$ ];
    count( $v$ ) = [ $c_0, \dots, c_{m_1+m_2}$ ] where
       $c_k = \sum_{i=0}^{m_1} \binom{k}{i} c_i^1 c_{k-i}^2$ ; //  $c_i^2 = 0$  for  $i \notin \{0, \dots, m_2\}$ 
    summands( $v$ ) = [ $s^0, \dots, s^{m_1+m_2}$ ] where  $s^k = [s_0^k, \dots, s_{m_1}^k]$  and
     $s_i^k = \begin{cases} \binom{k}{i} c_i^1 c_{k-i}^2 & \text{if } 0 \leq k-i \leq m_2 \\ 0 & \text{else} \end{cases}$ 
}

```

Figure 3.21: Algorithm Adorn

```

UnrankTreeNoCross( $r, v$ )
Input: a rank  $r$  and the root  $v$  of the SDG
Output: adorned SDG
let count( $v$ ) = [ $x_0, \dots, x_m$ ];
 $k := \min_j r \leq \sum_{i=0}^j x_i$ ; // efficiency: binary search on materialized sums.
 $r' := r - \sum_{i=0}^{k-1} x_i$ ;
UnrankLocalTreeNoCross( $v, r', k$ );

```

The following table shows the intervals associated with the partitions $\mathcal{T}_G^{e(k)}$ for the standard decomposition graph in Figure 3.20:

Partition	Interval
$\mathcal{T}_G^{e(1)}$	[1, 5]
$\mathcal{T}_G^{e(2)}$	[6, 10]
$\mathcal{T}_G^{e(3)}$	[11, 15]
$\mathcal{T}_G^{e(4)}$	[16, 18]

The unranking procedure makes use of unranking decompositions and unranking triplets. For the latter and a given X, Y, Z , we need to assign each member in

$$\{(x, y, z) | 1 \leq x \leq X, 1 \leq y \leq Y, 1 \leq z \leq Z\}$$

a unique number in $[1, XYZ]$ and base an unranking algorithm on this assignment. We leave this as a simple exercise to the reader and call the function `UnrankTriplet(r, X, Y, Z)`. Here, r is the rank and X, Y , and Z are the upper bounds for the numbers in the triplets. The code for unranking looks as follows:

```

UnrankingTreeNoCrossLocal( $v, r, k$ )
Input: an SDG node  $v$ , a rank  $r$ , a number  $k$  identifying a partition
Output: adornments of the SDG as a side-effect
Let  $\{w_1, \dots, w_n\}$  be the children of  $v$ 
switch ( $n$ ) {
  case 0:
    assert( $r = 1 \ \&\& \ k = 0$ );
    // no additional adornment for  $v$ 
  case 1:
    let  $\text{count}(v) = [c_0, \dots, c_n]$ ;
    let  $\text{summands}(v) = [s^0, \dots, s^n]$ ;
    assert( $k \leq n \ \&\& \ r \leq c_k$ );
     $k_1 = \min_j r \leq \sum_{i=0}^j s_i^k$ ;
     $r_1 = r - \sum_{i=0}^{k_1-1} s_i^k$ ;
    insert-at( $v$ ) =  $k$ ;
    UnrankingTreeNoCrossLocal( $w_1, r_1, k_1$ );
  case 2:
    let  $\text{count}(v) = [c_0, \dots, c_n]$ ;
    let  $\text{summands}(v) = [s^0, \dots, s^n]$ ;
    let  $\text{count}(w_1) = [c_0^1, \dots, c_{n_1}^1]$ ;
    let  $\text{count}(w_2) = [c_0^2, \dots, c_{n_2}^2]$ ;
    assert( $k \leq n \ \&\& \ r \leq c_k$ );
     $k_1 = \min_j r \leq \sum_{i=0}^j s_i^k$ ;
     $q = r - \sum_{i=0}^{k_1-1} s_i^k$ ;
     $k_2 = k - k_1$ ;
    ( $r_1, r_2, a$ ) = UnrankTriplet( $q, c_{k_1}^1, c_{k_2}^2, \binom{k}{i}$ );
     $\alpha = \text{UnrankDecomposition}(a)$ ;
    merge-using( $v$ ) =  $\alpha$ ;
    UnrankingTreeNoCrossLocal( $w_1, r_1, k_1$ );
    UnrankingTreeNoCrossLocal( $w_2, r_2, k_2$ );
}

```

3.3.4 Quick Pick

The QuickPick algorithm of Waas and Pellenkoft [784, 785] does not generate random join trees in the strong sense but comes close to it and is far easier to implement and more broadly applicable. The idea is to randomly select an edge in the query graph and to construct a join tree corresponding to this edge.

```

QuickPick(Query Graph  $G$ )
Input: a query graph  $G = (\{R_1, \dots, R_n\}, E)$ 
Output: a bushy join tree
BestTreeFound = any join tree
while stopping criterion not fulfilled {
     $E' = E$ ;
    Trees =  $\{R_1, \dots, R_n\}$ ;
    while ( $|\text{Trees}| > 1$ ) {
        choose  $e \in E'$ ;
         $E' - = e$ ;
        if ( $e$  connects two relations in different subtrees  $T_1, T_2 \in \text{Trees}$ ) {
            Trees  $- = T_1$ ;
            Trees  $- = T_2$ ;
            Trees  $+ = \text{CreateJoinTree}(T_1, T_2)$ ;
        }
    }
    Tree = single tree contained in Trees;
    if ( $\text{cost}(\text{Tree}) < \text{cost}(\text{BestTreeFound})$ ) {
        BestTreeFound = Tree;
    }
}
return BestTreeFound

```

3.3.5 Iterative Improvement

Swami and Gupta [747], Swami [746] and Ioannidis and Kang [385] applied the idea of iterative improvement to join ordering [385]. The idea is to start from a random plan and then to apply randomly selected transformations from a rule set if they improve the current join tree, until not further improvement is possible.

```

IterativeImprovementBase(Query Graph  $G$ )
Input: a query graph  $G = (\{R_1, \dots, R_n\}, E)$ 
Output: a join tree
do {
    JoinTree = random tree
    JoinTree = IterativeImprovement(JoinTree)
}

```

```

    if (cost(JoinTree) < cost(BestTree)) {
        BestTree = JoinTree;
    }
} while (time limit not exceeded)
return BestTree

```

IterativeImprovement(JoinTree)

Input: a join tree

Output: improved join tree

```

do {
    JoinTree' = randomly apply a transformation to JoinTree;
    if (cost(JoinTree') < cost(JoinTree)) {
        JoinTree = JoinTree';
    }
} while (local minimum not reached)
return JoinTree

```

The number of variants of iterative improvements is large. The first parameter is the used rule set. To restrict search to left-deep trees, a rule set consisting of *swap* and *3cycle* is appropriate [747]. If we consider bushy trees, a complete set consisting of commutativity, associativity, left join exchange and right join exchange makes sense. This rule set (proposed by Ioannidis and Kang) is appropriate to explore the whole space of bushy join trees. A second parameter is how to determine whether the local minimum has been reached. Considering all possible neighbor states of a join tree is expensive. Therefore, a subset of size k is sometimes considered. Then, for example, k can be limited to the number of edges in the query graph [747].

3.3.6 Simulated Annealing

Iterative Improvement suffers from the drawback that it only applies a move if it improves the current plan. This leads to the problem that one is often stuck in a local minimum. Simulated annealing tries to avoid this problem by allowing moves that result in more expensive plans [390, 385, 747]. However, instead of considering every plan, only those whose cost increase does not exceed a certain limit are considered. During time, this limit decreases. This general idea is cast into the notion temperatures and probabilities of performing a selected transformation. A generic formulation of simulated annealing could look as follows:

SimulatedAnnealing(Query Graph G)

Input: a query graph $G = (\{R_1, \dots, R_n\}, E)$

Output: a join tree

BestTreeSoFar = random tree;

Tree = BestTreeSoFar;

```

do {
  do {
    Tree' = apply random transformation to Tree;
    if (cost(Tree') < cost(Tree)) {
      Tree = Tree';
    } else {
      with probability  $e^{-(cost(Tree')-cost(Tree))/temperature}$ 
        Tree = Tree';
    }
    if (cost(Tree) < cost(BestTreeSoFar)) {
      BestTreeSoFar = Tree';
    }
  } while (equilibrium not reached)
reduce temperature;
} while (not frozen)
return BestTreeSoFar

```

Besides the rule set used, the initial temperature, the temperature reduction, and the definitions of equilibrium and frozen determine the algorithm's behavior. For each of them several alternatives have been proposed in the literature. The starting temperature can be calculated as follows: determine the standard deviation σ of costs by sampling and multiply it with a constant value ([747] use 20). An alternative is to set the starting temperature twice the cost of the first randomly selected join tree [385] or to determine the starting temperature such that at least 40% of all possible transformations are accepted [723].

For temperature reduction, we can apply the formula $temp* = 0.975$ [385] or $\max(0.5, e^{-\frac{\Delta t}{\sigma}})$ [747].

The equilibrium is defined to be reached if for example the cost distribution of the generated solutions is sufficiently stable [747], the number of iterations is sixteen times the number of relations in the query [385], or number of iterations is the same as the number of relations in the query [723].

We can establish frozenness if the difference between the maximum and minimum costs among all accepted join trees at the current temperature equals the maximum change in cost in any accepted move at the current temperature [747], the current solution could not be improved in four outer loop iterations and the temperature has been fallen below one [385], or the current solution could not be improved in five outer loop iterations and less than two percent of the generated moves were accepted [723].

Considering databases are used in mission critical applications. Would you bet your business on these numbers?

3.3.7 Tabu Search

Morzy, Matyasiak and Salza applied Tabu Search to join ordering [534]. The general idea is that among all neighbors reachable via the transformations, only

the cheapest is considered even if its cost are higher than the costs of the current join tree. In order to avoid running into cycles, a tabu set is maintained. It contains the last join trees generated, and the algorithm is not allowed to visit them again. This way, it can escape local minima, since eventually all nodes in the valley of a local minimum will be in the tabu set. The stopping conditions could be that there ws no improvement over the current best solution found during the last given number of iterations or if the set neighbors minus the tabu set is empty (in line (*)).

Tabu Search looks as follows:

```

TabuSearch(Query Graph)
Input: a query graph  $G = (\{R_1, \dots, R_n\}, E)$ 
Output: a join tree
Tree = random join tree;
BestTreeSoFar = Tree;
TabuSet =  $\emptyset$ ;
do {
  Neighbors = all trees generated by applying a transformation to Tree;
  Tree = cheapest in Neighbors  $\setminus$  TabuSet; (*)
  if (cost(Tree) < cost(BestTreeSoFar)) {
    BestTreeSoFar = Tree;
  }
  if(|TabuSet| > limit) remove oldest tree from TabuSet;
  TabuSet += Tree;
} while (not stopping condition satisfied);
return BestTreeSoFar;

```

3.3.8 Genetic Algorithms

Genetic algorithms are inspired by evolution: only the fittest survives [281]. They work with a population that evolves from generation to generation. Successors are generated by crossover and mutation. Further, a subset of the current population (the fittest) are propagated to the next generation (selection). The first generation is generated by a random generation process.

The problem is how to represent each individual in a population. The following analogies are used:

- Chromosome \longleftrightarrow string
- Gene \longleftrightarrow character

In order to solve an optimization problem with genetic algorithms, an encoding is needed as well as a specification for selection, crossover, and mutation.

Genetic algorithms for join ordering have been considered in [66, 723]. We first introduce alternative encodings, then come to the selection process, and finally discuss crossover and mutation.

Encodings We distinguish *ordered list* and *ordinal number* encodings. Both encodings are used for left-deep and bushy trees. In all cases we assume that the relations R_1, \dots, R_n are to be joined and use the index i to denote R_i .

1. Ordered List Encoding

(a) left-deep trees

A left-deep join tree is encoded by a permutation of $1, \dots, n$. For instance, $((R_1 \bowtie R_4) \bowtie R_2) \bowtie R_3$ is encoded as “1423”.

(b) bushy trees

Bennet, Ferris, and Ioannidis proposed the following encoding scheme [66, 67]. A bushy join-tree without cartesian products is encoded as an ordered list of the edges in the join graph. Therefore, we number the edges in the join graph. Then the join tree is encoded in a bottom-up, left-to-right manner. See Figure 3.22 for an example.

2. Ordinal Number Encoding

(a) left-deep trees

A join tree is encoded by using a list of relations that is shortened whenever a join has been encoded. We start with the list $L = \langle R_1, \dots, R_n \rangle$. Then within L we find the index of first relation to be joined. Let this relation be R_i . R_i is the i -th relation in L . Hence, the first character in the chromosome string is i . We eliminate R_i from L . For every subsequent relation joined, we again determine its index in L , remove it from L and append the index to the chromosome string. For instance, starting with $\langle R_1, R_2, R_3, R_4 \rangle$, the left-deep join tree $((R_1 \bowtie R_4) \bowtie R_2) \bowtie R_3$ is encoded as “1311”.

(b) bushy trees

Again, we start with the list $L = \langle R_1, \dots, R_n \rangle$ and encode a bushy join tree in a bottom-up, left-to-right manner. Let $R_i \bowtie R_j$ be the first join in the join tree under this ordering. Then we look up their positions in L and add them to the encoding. Next we eliminate R_i and R_j from L and push $R_{i,j}$ to the front of it. We then proceed for the other joins by again selecting the next join which now can be between relations and/or subtrees. We determine their position within L , add these positions to the encoding, remove them from L , and insert a composite relation into L such that the new composite relation directly follows those already present. For instance, starting with the list $\langle R_1, R_2, R_3, R_4 \rangle$, the bushy join tree $((R_1 \bowtie R_2) \bowtie (R_3 \bowtie R_4))$ is encoded as “12 23 12”.

The encoding is completed by adding join methods.

Crossover A crossover generates a new solution from two individuals. Therefore, two partial solutions are combined. Obviously, its definition depends on the encoding. Two kinds of crossovers are distinguished: the subsequence and the subset exchange.




Figure 3.22: A query graph, a join tree, and its encoding

The subsequence exchange for the ordered list encoding works as follows. Assume two individuals with chromosomes $u_1v_1w_1$ and $u_2v_2w_2$. From these we generate $u_1v'_1w_1$ and $u_2v'_2w_2$, where v'_i is a permutation of the relations in v_i such that the order of their appearance is the same as in $u_{3-i}v_{3-i}w_{3-i}$. In order to adapt the subsequence exchange operator to the ordinal number encoding, we have to require that the v_i are of equal length ($|v_1| = |v_2|$) and occur at the same offset ($|u_1| = |u_2|$). We then simply swap the v_i . That is, we generate $u_1v_2w_1$ and $u_2v_1w_2$.

The subset exchange is defined only for the ordered list encoding. Within the two chromosomes, we find two subsequences of equal length comprising the same set of relations. These sequences are then simply exchanged.

Mutation A mutation randomly alters a character in the encoding. If duplicates must not occur — as in the ordered list encoding — swapping two characters is a perfect mutation.

Selection The probability of a join tree's survival is determined by its rank in the population. That is, we calculate the costs of the join trees encoded for each member of the population. Then we sort the population according to their associated costs and assign probabilities to each individual such that the best solution in the population has the highest probability to survive and so on. After probabilities have been assigned, we randomly select members of the population taking these probabilities into account. That is, the higher the probability of a member, the higher is its chance to survive.

Algorithm The genetic algorithm then works as follows. First, we create a random population of a given size (say 128). We apply crossover and mutation with a given rate, for example such that 65% of all members of a population participate in crossover, and 5% of all members of a population are subject to random mutation. Then we apply selection until we again have a population of a given size. We stop after we have not seen an improvement within the

population for a fixed number of iterations (say 30).

3.4 Hybrid Algorithms

All the algorithms we have seen so far can be combined to result in new approaches to join ordering. Some of the numerous possibilities have been described in the literature. We present them.

3.4.1 Two Phase Optimization

Two phase optimization combines Iterative Improvement with Simulated Annealing [385]. For a number of randomly generated initial trees, Iterative Improvement is used to find a local minimum. Then Simulated Annealing is started to find a better plan in the neighborhood of the local minima. The initial temperature of Simulated Annealing can be lower as is its original variants.

3.4.2 AB-Algorithm

The AB-Algorithm was developed by Swami and Iyer [748, 749]. It builds on the IKKBZ-Algorithm by resolving its limitations. First, if the query graph is cyclic, a spanning tree is selected. Second, two different cost functions for joins (join methods) are supported by the AB-Algorithm: nested loop join and sort merge join. In order to make the sort merge join's cost model fit the ASI property, it is simplified. Third, join methods are assigned randomly before IKKBZ is called. Afterwards, an iterative improvement phase follows. The algorithm can be formulated as follows:

```

AB(Query Graph  $G$ )
Input: a query graph  $G = (\{R_1, \dots, R_n\}, E)$ 
Output: a left-deep join tree
while (number of iterations  $\leq n^2$ ) {
  if  $G$  is cyclic take spanning tree of  $G$ 
  randomly attach a join method to each relation
  JoinTree = result of IKKBZ
  while (number of iterations  $\leq n^2$ ) {
    apply Iterative Improvement to JoinTree
  }
}
return best tree found

```

3.4.3 Toured Simulated Annealing

Lanzelotte, Valduriez, and Zäit introduced *toured simulated annealing* as a search strategy useful in distributed databases where the search space is even

larger than in centralized systems [454]. The basic idea is that simulated annealing is called n times with different initial join trees, if n is the number of relations to be joined. Each join sequence in the set `Solutions` produced by `GreedyJoinOrdering-3` is used to start an independent run of simulated annealing. As a result, the starting temperature can be decreased to 0.1 times the cost of the initial plan.

3.4.4 GOO-II

GOO-II appends an Iterative Improvement step to the GOO-Algorithm.

3.4.5 Iterative Dynamic Programming

Iterative Dynamic Programming combines heuristics with dynamic programming in order to overcome the deficiencies of both. It comes in two variants [440, 696]. The first variant, `IDP-1` (see Figure 3.23), first creates all join trees which contain up to k relations where k is a parameter of the algorithm. After this step, it selects the cheapest join tree comprising k relations, replaces it by a new compound relation and starts all over again. The iteration stops, when only one compound relation representing a join tree for all relations remains in the `ToDo` list.

The second variant, `IDP-2` (see Figure 3.24), works the other way round. It first applies a greedy heuristics to build join trees of size up to k . To the larger subtree it applies dynamic programming to improve it. The result of the optimized outcome of the greedy algorithm is then encapsulated in a new compound relation which replaces its constituent relations in the `ToDo` list. The algorithm then iterates until only one entry remains in the `ToDo` list.

Obviously, from these two basic variants several others can be derived. A systematic investigation of the basic algorithms and their variants is given by Kossmann and Stocker [440]. It turns out that the most promising variants exist for `IDP-1`.

3.5 Ordering Order-Preserving Joins

This section covers an algorithm for ordering order-preserving joins [528]. This is important for XQuery and other languages that require order-preservation. XQuery specifies that the result of a query is a sequence. If no `unordered` or `order by` instruction is given, the order of the output sequence is determined by the order of the input sequences given in the `for` clauses of the query. If there are several entries in a `for` clause or several `for` clauses, order-preserving join operators [155] can be a natural component for the evaluation of such a query.

The order-preserving join operator is used in several algebras in the context of

- semi-structured data and XML (e.g. SAL [63], XAL [249]),
- OLAP [711], and

```

IDP-1( $\{R_1, \dots, R_n\}$ ,  $k$ )
Input: a set of relations to be joined, maximum block size  $k$ 
Output: a join tree
for ( $i = 1$ ;  $i \leq n$ ;  $++i$ ) {
    BestTree( $\{R_i\}$ ) =  $R_i$ ;
}
ToDo =  $\{R_1, \dots, R_n\}$ ;
while ( $|ToDo| > 1$ ) {
     $k = \min(k, |ToDo|)$ ;
    for ( $i = 2$ ;  $i < k$ ;  $++i$ ) {
        for all  $S \subseteq ToDo$ ,  $|S| = i$  do {
            for all  $O \subset S$  do {
                BestTree( $S$ ) = CreateJoinTree(BestTree( $S \setminus O$ ), BestTree( $O$ ));
            }
        }
    }
    find  $V \subset ToDo$ ,  $|V| = k$ 
    with  $\text{cost}(\text{BestTree}(V)) = \min\{\text{cost}(\text{BestTree}(W)) \mid W \subset ToDo, |W| = k\}$ ;
    generate new symbol  $T$ ;
    BestTree( $\{T\}$ ) = BestTree( $V$ );
    ToDo =  $(ToDo \setminus V) \cup \{T\}$ ;
    for all  $O \subset V$  do delete(BestTree( $O$ ));
}
return BestTree( $\{R_1, \dots, R_n\}$ );

```

Figure 3.23: Pseudo code for IDP-1

- time series data [466].

We give a polynomial algorithm that produces bushy trees for a sequence of order-preserving joins and selections. These trees may contain cross products even if the join graph is connected. However, we apply selections as early as possible. The algorithm then produces the optimal plan among those who push selections down. The cost function is a parameter of the algorithm, and we do not need to restrict ourselves to those having the ASI property. Further, we need no restriction on the join graph, i.e. the algorithm produces the optimal plan even if the join graph is cyclic.

Before defining the order-preserving join, we need some preliminaries. The above algebras work on sequences of sets of variable bindings, i.e. sequences of unordered tuples where every attribute corresponds to a variable. (See Chapter 7.3 for a general discussion.) Single tuples are constructed using the standard $[\cdot]$ brackets. Concatenation of tuples and functions is denoted by \circ . The set of attributes defined for an expression e is defined as $\mathcal{A}(e)$. The set of free variables of an expression e is defined as $\mathcal{F}(e)$. For sequences e , we use

```

IDP-2( $\{R_1, \dots, R_n\}, k$ )
Input: a set of relations to be joined, maximum block size  $k$ 
Output: a join tree
for ( $i = 1; i \leq n; ++i$ ) {
    BestTree( $\{R_i\}$ ) =  $R_i$ ;
}
ToDo =  $\{R_1, \dots, R_n\}$ ;
while ( $|\text{ToDo}| > 1$ ) {
    // apply greedy algorithm to select a good building block
     $B = \emptyset$ ;
    for all  $v \in \text{ToDo}$ , do {
         $B += \text{BestTree}(\{v\})$ ;
    }
    do {
        find  $L, R \in B$ 
        with  $\text{cost}(\text{CreateJoinTree}(L, R))$ 
            =  $\min\{\text{cost}(\text{CreateJoinTree}(L', R')) \mid L', R' \in B\}$ ;
         $P = \text{CreateJoinTree}(L, R)$ ;
         $B = (B \setminus \{L, R\}) \cup \{P\}$ ;
    } while ( $P$  involves no more than  $k$  relations and  $|B| > 1$ );
    // reoptimize the bigger of  $L$  and  $R$ ,
    // selected in the last iteration of the greedy loop
    if ( $L$  involves more tables than  $R$ ) {
        ReOpRels = relations involved in  $L$ ;
    } else {
        ReOpRels = relations involved in  $R$ ;
    }
     $P = \text{DP-Bushy}(\text{ReOpRels})$ ;
    generate new symbol  $T$ ;
    BestTree( $\{T\}$ ) =  $P$ ;
    ToDo =  $(\text{ToDo} \setminus \text{ReOpRels}) \cup \{T\}$ ;
    for all  $O \subset V$  do delete(BestTree( $O$ ));
}
return BestTree( $\{R_1, \dots, R_n\}$ );

```

Figure 3.24: Pseudocode for IDP-2

$\alpha(e)$ to denote the first element of a sequence. We identify single element sequences with elements. The function τ retrieves the tail of a sequence, and \oplus concatenates two sequences. We denote the empty sequence by ϵ .

We define the algebraic operators recursively on their input sequences. The order-preserving join operator is defined as the concatenation of an order-preserving selection and an order-preserving cross product. For unary operators, if the input sequence is empty, the output sequence is also empty. For

binary operators, the output sequence is empty whenever the left operand represents an empty sequence.

The order-preserving join operator is based on the definition of an order-preserving cross product operator defined as

$$e_1 \hat{\times} e_2 := (\alpha(e_1) \hat{\times} e_2) \oplus (\tau(e_1) \hat{\times} e_2)$$

where

$$e_1 \hat{\times} e_2 := \begin{cases} \epsilon & \text{if } e_2 = \epsilon \\ (e_1 \circ \alpha(e_2)) \oplus (e_1 \hat{\times} \tau(e_2)) & \text{else} \end{cases}$$

We are now prepared to define the join operation on ordered sequences:

$$e_1 \bowtie_p e_2 := \hat{\sigma}_p(e_1 \hat{\times} e_2)$$

where the order-preserving selection is defined as

$$\hat{\sigma}_p(e) := \begin{cases} \epsilon & \text{if } e = \epsilon \\ \alpha(e) \oplus \hat{\sigma}_p(\tau(e)) & \text{if } p(\alpha(e)) \\ \hat{\sigma}_p(\tau(e)) & \text{else} \end{cases}$$

As usual, selections can be reordered and pushed inside order-preserving joins. Besides, the latter are associative. The following equivalences formalize this.

$$\begin{aligned} \hat{\sigma}_{p_1}(\hat{\sigma}_{p_2}(e)) &= \hat{\sigma}_{p_2}(\hat{\sigma}_{p_1}(e)) \\ \hat{\sigma}_{p_1}(e_1 \bowtie_{p_2} e_2) &= \hat{\sigma}_{p_1}(e_1) \bowtie_{p_2} e_2 & \text{if } \mathcal{F}(p_1) \subseteq \mathcal{A}(e_1) \\ \hat{\sigma}_{p_1}(e_1 \bowtie_{p_2} e_2) &= e_1 \bowtie_{p_2} \hat{\sigma}_{p_1}(e_2) & \text{if } \mathcal{F}(p_1) \subseteq \mathcal{A}(e_2) \\ e_1 \bowtie_{p_1}(e_2 \bowtie_{p_2} e_3) &= (e_1 \bowtie_{p_1} e_2) \bowtie_{p_2} e_3 & \text{if } \mathcal{F}(p_i) \subseteq \mathcal{A}(e_i) \cup \mathcal{A}(e_{i+1}) \end{aligned}$$

While being associative, the order-preserving join is not commutative, as the following example illustrates. Given two tuple sequences $R_1 = \langle [a : 1], [a : 2] \rangle$ and $R_2 = \langle [b : 1], [b : 2] \rangle$, we have

$$\begin{aligned} R_1 \bowtie_{true} R_2 &= \langle [a : 1, b : 1], [a : 1, b : 2], [a : 2, b : 1], [a : 2, b : 2] \rangle \\ R_2 \bowtie_{true} R_1 &= \langle [a : 1, b : 1], [a : 2, b : 1], [a : 1, b : 2], [a : 2, b : 2] \rangle \end{aligned}$$

Before introducing the algorithm, let us have a look at the size of the search space. Since the order-preserving join is associative but not commutative, the input to the algorithm must be a sequence of join operators or, likewise, a sequence of relations to be joined. The output is then a fully parenthesized expression. Given a sequence of n binary associative but not commutative operators, the number of fully parenthesized expressions is (see [178])

$$P(n) = \begin{cases} 1 & \text{if } n = 1 \\ \sum_{k=1}^{n-1} P(k)P(n-k) & \text{if } n > 1 \end{cases}$$

We have that $P(n) = C(n-1)$, where $C(n)$ are the Catalan numbers defined as $C(n) = \frac{1}{n+1} \binom{2n}{n}$. Since $C(n) = \Omega(\frac{4^n}{n^{3/2}})$, the search space is exponential in size.

```

applicable-predicates( $\mathcal{R}$ ,  $\mathcal{P}$ )
01   $\mathcal{B} = \emptyset$ 
02  foreach  $p \in \mathcal{P}$ 
03      IF ( $\mathcal{F}(p) \subseteq \mathcal{A}(\mathcal{R})$ )
04           $\mathcal{B}+ = p$ 
05  return  $\mathcal{B}$ 

```

Figure 3.25: Subroutine `applicable-predicates`

The algorithm is inspired by the dynamic programming algorithm for finding optimal parenthesized expressions for matrix-chain multiplication [178]. The differences are that we have to encapsulate the cost function and deal with selections. We give a detailed example application of the algorithm below. This example illustrates (1) the optimization potential, (2) that cross products can be favorable, (3) how to plug in a cost function into the algorithm, and (4) the algorithm itself.

The algorithm itself is broken up into several subroutines. The first is `applicable-predicates` (see Fig. 3.25). Given a sequence of relations R_i, \dots, R_j and a set of predicates, it retrieves those predicates applicable to the result of the join of the relations. Since joins and selections can be reordered freely, the only condition for a predicate to be applicable is that all its free variables are bound by the given relations.

The second subroutine is the most important and intriguing. It fills several arrays with values in a bottom-up manner. The third subroutine then builds the query evaluation plan using the data in the arrays.

The subroutine `construct-bushy-tree` takes as input a sequence R_1, \dots, R_n of relations to be joined and a set \mathcal{P} of predicates to be applied. For every possible subsequence R_i, \dots, R_j , the algorithm finds the best plan to join these relations. Therefore, it determines some k such that the cheapest plan joins the intermediate results for R_i, \dots, R_k and R_{k+1}, \dots, R_j by its topmost join. For this it is assumed that for all k the best plans for joining R_i, \dots, R_k and R_{k+1}, \dots, R_j are known. Instead of directly storing the best plan, we remember (1) the costs of the best plan for R_i, \dots, R_j for all $1 \leq i \leq j \leq n$ and (2) the k where the split takes place. More specifically, the array $c[i, j]$ contains the costs of the best plan for joining R_i, \dots, R_j , and the array $t[i, j]$ contains the k such that this best plan joins R_i, \dots, R_k and R_{k+1}, \dots, R_j with its topmost join. For every sequence R_i, \dots, R_j , we also remember the set of predicates that can be applied to it, excluding those that have been applied earlier. These applicable predicates are contained in $p[i, j]$. Still, we are not done. All cost functions we know use some kind of statistics on the argument relation(s) in order to compute the costs of some operation. Since we want to be generic with respect to the cost function, we encapsulate the computation of statistics and costs within functions S_0 , C_0 , S_1 , and C_1 . The function S_0 retrieves statistics for base relations. The function C_0 computes the costs of retrieving (part of) a base relation. Both functions take a set of applicable predicates as an additional


```

construct-bushy-tree( $\mathcal{R}$ ,  $\mathcal{P}$ )
01   $n = |\mathcal{R}|$ 
02  for  $i = 1$  to  $n$ 
03     $\mathcal{B} = \text{applicable-predicates}(R_i, \mathcal{P})$ 
04     $\mathcal{P} = \mathcal{P} \setminus \mathcal{B}$ 
05     $p[i, i] = \mathcal{B}$ 
06     $s[i, i] = S_0(R_i, \mathcal{B})$ 
07     $c[i, i] = C_0(R_i, \mathcal{B})$ 
08  for  $l = 2$  to  $n$ 
09    for  $i = 1$  to  $n - l + 1$ 
10       $j = i + l - 1$ 
11       $\mathcal{B} = \text{applicable-predicates}(R_{i\dots j}, \mathcal{P})$ 
12       $\mathcal{P} = \mathcal{P} \setminus \mathcal{B}$ 
13       $p[i, j] = \mathcal{B}$ 
14       $s[i, j] = S_1(s[i, j - 1], s[j, j], \mathcal{B})$ 
15       $c[i, j] = \infty$ 
16      for  $k = i$  to  $j - 1$ 
17         $q = c[i, k] + c[k + 1, j] + C_1(s[i, k], s[k + 1, j], \mathcal{B})$ 
18        IF ( $q < c[i, j]$ )
19           $c[i, j] = q$ 
20           $t[i, j] = k$ 

```

Figure 3.26: Subroutine `construct-bushy-tree`

```

extract-plan( $\mathcal{R}$ ,  $t$ ,  $p$ )
01  return extract-subplan( $\mathcal{R}$ ,  $t$ ,  $p$ , 1,  $|\mathcal{R}|$ )

extract-subplan( $\mathcal{R}$ ,  $t$ ,  $p$ ,  $i$ ,  $j$ )
01  IF ( $j > i$ )
02     $X = \text{extract-subplan}(\mathcal{R}, t, p, i, t[i, j])$ 
03     $Y = \text{extract-subplan}(\mathcal{R}, t, p, t[i, j] + 1, j)$ 
04    return  $X \bowtie_{p[i, j]} Y$ 
05  else
06    return  $\hat{\sigma}_{p[i, i]}(R_i)$ 

```

Figure 3.27: Subroutine `extract-plan` and its subroutine

argument. The function S_1 computes the statistics for intermediate relations. Since the result of joining some relations R_i, \dots, R_j may occur in many different plans, we compute it only once and store it in the array s . C_1 computes the costs of joining two relations and applying a set of predicates. Below, we show how concrete (simple) cost and statistics functions can look like.

Given the above, the algorithm (see Fig. 3.26) fills the arrays in a bottom-up manner by first computing for every base relation the applicable predicates, the

statistics of the result of applying the predicates to the base relation and the costs for computing these intermediate results, i.e. for retrieving the relevant part of the base relation and applying the predicates (lines 02-07). Note that this is not really trivial if there are several index structures that can be applied. Then computing C_0 involves considering different access paths. Since this is an issue orthogonal to join ordering, we do not detail on it.

After we have the costs and statistics for sequences of length one, we compute the same information for sequences of length two, three, and so on until n (loop starting at line 08). For every length, we iterate over all subsequences of that length (loop starting at line 09). We compute the applicable predicates and the statistics. In order to determine the minimal costs, we have to consider every possible split point. This is done by iterating the split point k from i to $j - 1$ (line 16). For every k , we compute the cost and remember the k that resulted in the lowest costs (lines 17-20).

The last subroutine takes the relations, the split points (t), and the applicable predicates (p) as its input and extracts the plan. The whole plan is extracted by calling `extract-plan`. This is done by instructing `extract-subplan` to retrieve the plan for all relations. This subroutine first determines whether the plan for a base relation or that of an intermediate result is to be constructed. In both cases, we did a little cheating here to keep things simple. The plan we construct for base relations does not take the above-mentioned index structures into account but simply applies a selection to a base relation instead. Obviously, this can easily be corrected. We also give the join operator the whole set of predicates that can be applied. That is, we do not distinguish between join predicates and other predicates that are better suited for a selection subsequently applied to a join. Again, this can easily be corrected.

Let us have a quick look at the complexity of the algorithm. Given n relations with m attributes in total and p predicates, we can implement `applicable-predicates` in $O(pm)$ by using a bit vector representation for attributes and free variables and computing the attributes for each sequence R_i, \dots, R_j once upfront. The latter takes $O(n^2m)$.

The complexity of the routine `construct-bushy-tree` is determined by the three nested loops. We assume that S_1 and C_1 can be computed in $O(p)$, which is quite reasonable. Then, we have $O(n^3p)$ for the innermost loop, $O(n^2)$ calls to `applicable-predicates`, which amounts to $O(n^2pm)$, and $O(n^2p)$ for calls of S_1 . Extracting the plan is linear in n . Hence, the total runtime of the algorithm is $O(n^2(n + m)p)$

In order to illustrate the algorithm, we need to fix the functions S_0 , S_1 , C_0 and C_1 . We use the simple cost function C_{out} . As a consequence, the array s simply stores cardinalities, and S_0 has to extract the cardinality of a given base relation and multiply it by the selectivities of the applicable predicates. S_1 multiplies the input cardinalities with the selectivities of the applicable predicates. We set C_0 to zero and C_1 to S_1 . The former is justified by the fact that every relation must be accessed exactly once and hence, the access costs are equal in

all plans. Summarizing, we define

$$\begin{aligned}
 S_0(R, \mathcal{B}) &:= |R| \prod_{p \in \mathcal{B}} f(p) \\
 S_1(x, y, \mathcal{B}) &:= xy \prod_{p \in \mathcal{B}} f(p) \\
 C_0(R, \mathcal{B}) &:= 0 \\
 C_1(x, y, \mathcal{B}) &:= S_1(x, y, \mathcal{B})
 \end{aligned}$$

where \mathcal{B} is a set of applicable predicates and for a single predicate p , $f(p)$ returns its selectivity.

We illustrate the algorithm by an example consisting of four relations R_1, \dots, R_4 with cardinalities $|R_1| = 200$, $|R_2| = 1$, $|R_3| = 1$, $|R_4| = 20$. Besides, we have three predicates $p_{i,j}$ with $\mathcal{F}(p_{i,j}) \subseteq \mathcal{A}(R_i) \cup \mathcal{A}(R_j)$. They are $p_{1,2}$, $p_{3,4}$, and $p_{1,4}$ with selectivities $1/2$, $1/10$, $1/5$.

Let us first consider an example plan and its costs. The plan

$$((R_1 \bowtie_{p_{1,2}} R_2) \bowtie_{true} R_3) \bowtie_{p_{1,4} \wedge p_{3,4}} R_4$$

has the costs $240 = 100 + 100 + 40$.

For our simple cost function, the algorithm `construct-bushy-tree` will fill the array s with the initial values:

s			
200			
	1		
		1	
			20

After initialization, the array c has 0 everywhere in its diagonal and the array p empty sets.

For $l = 2$, the algorithm produces the following values:

l	i	j	k	s[i,j]	q	current c[i,j]	current t[i,j]
2	1	2	1	100	100	100	1
2	2	3	2	1	1	1	2
2	3	4	3	2	2	2	3

For $l = 3$, the algorithm produces the following values:

l	i	j	k	s[i,j]	q	current c[i,j]	current t[i,j]
3	1	3	1	200	101	101	1
3	1	3	2	200	200	101	1
3	2	4	2	2	4	4	2
3	2	4	3	2	3	3	3

For $l = 4$, the algorithm produces the following values:

l	i	j	k	s[1,4]	q	current c[1,4]	current t[1,4]
4	1	4	1	40	43	43	1
4	1	4	2	40	142	43	1
4	1	4	3	40	141	43	1

where for each k the value of q (in the following table denoted by q_k) is determined as follows:

$$\begin{aligned} q_1 &= c[1, 1] + c[2, 4] + 40 = 0 + 3 + 40 = 43 \\ q_2 &= c[1, 2] + c[3, 4] + 40 = 100 + 2 + 40 = 142 \\ q_3 &= c[1, 3] + c[4, 4] + 40 = 101 + 0 + 40 = 141 \end{aligned}$$

Collecting all the above $t[i, j]$ values leaves us with the following array as input for `extract-plan`:

$i \setminus j$	1	2	3	4
1		1	1	1
2			2	3
3				3
4				

The function `extract-plan` merely calls `extract-subplan`. For the latter, we give the call hierarchy and the result produced:

```

000 extract-plan(..., 1, 4)
100   extract-plan(..., 1, 1)
200   extract-plan(..., 2, 4)
210     extract-plan(..., 2, 3)
211       extract-plan(..., 2, 2)
212       extract-plan(..., 3, 3)
210     return (R2⋈trueR3)
220     extract-plan(..., 4, 4)
200   return ((R2⋈trueR3)⋈p3,4R4)
000 return (R1⋈p1,2∧p1,4((R2⋈trueR3)⋈p3,4R4))

```

The total cost of this plan is $c[1, 4] = 43$.

3.6 Characterizing Search Spaces

3.6.1 Complexity Thresholds

The complexity results presented in Section 3.1.6 show that most classes of join ordering problems are NP-hard. However, it is quite clear that some instances of the join ordering problem are simpler than others. For example, consider a query graph which is a clique in n relations R_1, \dots, R_n . Further assume that each R_i has cardinality 2^i and all join selectivities are $1/2$ (i.e. $f_{i,j} = 1/2$ for all $1 \leq i, j \leq n$, $i \neq j$). Obviously, this problem is easy to optimize although the query graph is clique. In this section we present some ideas on how the complexity of instances of the join ordering problem is influenced by certain parameters.

How can we judge the complexity of a single instance of a join ordering problem? Using standard complexity theory, for single problem instances we easily derive an algorithm that works in $\Theta(1)$. Hence, we must define other complexity measures. Consider our introductory join ordering problem. A simple greedy algorithm that orders relations according to their cardinality produces an optimal solution for it. Hence, one possibility to define the problem complexity would be how far a solution produced by typical heuristics for join ordering differ from the optimal solution. Another possibility is to use randomized algorithms like iterative improvement of simulated annealing and see how far the plans generated by them deviate from the optimal plan. These approaches have the problem that the results may depend on the chosen algorithm. This can be avoided by using the following approach. For each join ordering problem instance, we compute the fraction of good plans compared to all plans. Therefore, we need a measure of “good”. Typical examples thereof would be to say a plan is “good” if it does not deviate more than 10% or a factor of two from the optimal plan.

If these investigations were readily available, there are certain obvious benefits [438]:

1. The designer of an optimizer can classify queries such that heuristics are applied where they guarantee success; cases where they are bound to fail can be avoided. Furthermore, taking into account the vastly different run time of the different join ordering heuristics and probabilistic optimization procedures, the designer of an optimizer can choose the method that achieves a satisfactory result with the least effort.
2. The developer of search procedures and heuristics can use this knowledge to design methods solving hard problems (as exemplified for graph coloring problems [370]).
3. The investigator of different join ordering techniques is able to (1) consciously design challenging benchmarks and (2) evaluate existing benchmarks according to their degree of challenge.

The kind of investigation presented in this section first started in the context of artificial intelligence where a paper by Cheeseman, Kanefsky, and Taylor [138] spurred a whole new branch of research where the measures to judge the complexity of problem instances was investigated for many different NP-complete problems like satisfiability [138, 181, 277, 521], graph coloring [138], Hamiltonian circuits [138], traveling salesman [138], and constraint satisfaction [805].

We only present a small fraction of all possible investigations. The restrictions are that we do not consider all parameters that possibly influence the problem complexity, we only consider left-deep trees, and we restrict ourselves to the cost function C_{hj} . The join graphs are randomly generated. Starting with a circle, we randomly added edges until a clique is reached. The reader is advised to carry out his or her own experiments. Therefore, the following pointer into the literature might be useful. Lanzelotte and Valduriez provide an

object-oriented design for search strategies [452]. This allows easy modification and even the exchange of the plan generator's search strategy.

Search Space Analysis

The goal of this section is to determine the influence of the parameters on the search space of left-deep join trees. More specifically, we are interested in how a variation of the parameters changes the percentage of good solutions among all solutions. The quality of a solution is measured by the factor its cost deviates from the optimal permutation. For this, all permutations have to be generated and evaluated. The results of this experiment are shown in Figures 3.28 and 3.29. Each single curve accumulates the percentage of all permutations deviating less than a certain factor (given as the label) from the optimum. The accumulated percentages are given at the y-axes, the connectivity at the x-axes. The connectivity is given by the number of edges in the join graph. The curves within the figures are organized as follows. Figure 3.28 (3.29) shows varying mean selectivity values (relation sizes) and variances where the mean selectivity values (relation sizes) increase from top to bottom and the variances increase from left to right.

Note that the more curves are visible and the lower their y-values, the harder is the problem. We observe the following:

- all curves exhibit a minimum value at a certain connectivity
- which moves with increasing mean values to the right;
- increasing variances does not have an impact on the *minimum connectivity*,
- problems become less difficult with increasing mean values.

These findings can be explained as follows. With increasing connectivity, the join ordering problem becomes more complex up to a certain point and then less complex again. To see this, consider the following special though illustrative case. Assume an almost equal distribution of the costs of all alternatives between the worst case and optimal costs, equal relation sizes, and equal selectivities. Then the optimization potential *worst case/optimum* is 1 for connectivity 0 and cliques. In between, there exists a connectivity exhibiting the maximum optimization potential. This connectivity corresponds to the minimum connectivity of Figures 3.28 and 3.29.

There is another factor which influences the complexity of a single problem instance. Consider joining n relations. The problem becomes less complex if after joining $i < n$ relations the intermediate result becomes so small that the accumulated costs of the subsequent $n - i$ joins are small compared to the costs of joining the first i relations. Hence, the ordering of the remaining $n - i$ relations does not have a big influence on the total costs. This is the case for very small relations, small selectivities, or high connectivities. The greater selectivities and relation sizes are, the more relations have to be joined to reach this critical size of the intermediate result. If the connectivity is enlarged, this critical size is reached earlier. Since the number of selectivities involved in the

first few joins is small regardless of the connectivity, there is a lower limit to the number of joined relations required to arrive at the critical intermediate result size. If the connectivity is larger, this point is reached earlier, but there exists a lower limit on the connectivity where this point is reached. The reason for this lower limit is that the number of selectivities involved in the joins remains small for the first couple of relations, independent of their connectivity. These lines of argument explain subsequent findings, too.

The reader should be aware of the fact that the number of relations joined is quite small (10) in our experiments. Further, as observed by several researchers, if the number of joins increases, the number of “good” plans decreases [257, 745]. That is, increasing the number of relations makes the join ordering problem more difficult.

Figure 3.28: Impact of selectivity on the search space

Figure 3.29: Impact of relation sizes on the search space

Heuristics

For analyzing the influence of the parameters on the performance of heuristics, we give the figures for four different heuristics. The first two are very simple. The *minSel* heuristic selects those relations first of which incident join edges exhibit the minimal selectivity. The *recMinRel* heuristic chooses those relations first which result in the smallest intermediate relation.

We also analyzed the two advanced heuristics *IKKBZ* and *RDC*. The *IKKBZ* heuristic [443] is based on an optimal join ordering procedure [377, 443] which is applied to the minimal spanning tree of the join graph where the edges are labeled by the selectivities. The family of *RDC* heuristics is based on the relational difference calculus as developed in [358]. Since our goal is not to benchmark different heuristics in order to determine the best one, we have chosen the simplest variant of the family of *RDC* based heuristics. Here, the relations are ordered according to a certain weight whose actual computation is—for the purpose of this section—of no interest. The results of the experiments are presented in Figure 3.30.

On a first glance, these figures look less regular than those presented so far. This might be due to the non-stable behavior of the heuristics. Nevertheless, we can extract the following observations. Many curves exhibit a peak at a certain connectivity. Here, the heuristics perform worst. The peak connectivity is dependent on the selectivity size but not as regular as in the previous curves. Further, higher selectivities flatten the curves, that is, heuristics perform better at higher selectivities.

Figure 3.30: Impact of parameters on the performance of heuristics

Probabilistic Optimization Procedures

Figure 3.31 shows four pictures corresponding to simulated annealing (SA), iterative improvement (II), iterative improvement applied to the outcome of the IKKBZ heuristic (IKKBZ/II) and the RDC heuristic (RDC/II) [358]. The patterns shown in Figure 3.31 are very regular. All curves exhibit a peak at a certain connectivity. The peak connectivities typically coincide with the minimum connectivity of the search space analysis. Higher selectivities result in flatter curves; the probabilistic procedures perform better. These findings are absolutely coherent with the search space analysis. This is not surprising, since the probabilistic procedures investigate systematically —although with some random influence— a certain part of the search space.

Given a join ordering problem, we can describe its potential search space as a graph. The set of nodes consists of the set of join trees. For every two join trees a and b , we add an edge (a, b) if b can be reached from a by one of the transformation rules used in the probabilistic procedure. Further, with every node we can associate the cost its corresponding join tree.

Having in mind that the probabilistic algorithms are always in danger of being stuck in a local minima, the following two properties of the search space are of interest:

1. the cost distribution of local minima, and
2. the connection cost of low local minima.

Of course, if all local minima are of about the same cost, we do not have to worry, otherwise we do. It would be very interesting to know the percentage of local minima that are close to the global minima.

Concerning the second property, we first have to define the connection cost. Let a and b be two nodes and P be the set of all paths from a to b . The *connection cost* of a and b is then defined as $\min_{p \in P} \max_{s \in p} \{cost(s) | s \neq a, s \neq b\}$. Now, if the connection costs are high, we know that if we have to travel from one local minima to another, there is at least one node we have to pass which has high costs. Obviously, this is bad for our probabilistic procedures. Ioannidis and Kang [386] call a search graph that is favorable with respect to the two properties a *well*. Unfortunately, investigating these two properties of real search spaces is rather difficult. However, Ioannidis and Kang, later supported by Zhang, succeeded in characterizing cost wells in random graphs [386, 387]. They also conclude that the search space comprising bushy trees is better w.r.t. our two properties than the one for left-deep trees.

Figure 3.31: Impact of selectivities on probabilistic procedures

3.7 Discussion

Choose one of dynamic programming, memoization, permutations as the core of your plan generation algorithm and extend it with the rest of book.

ToDo

3.8 Bibliography

ToDo: Oezsu, Meechan [563, 564]

Chapter 4

Database Items, Building Blocks, and Access Paths

In this chapter we go down to the storage layer and discuss leaf nodes of query execution plans and plan fragments. We briefly recap some notions, but reading a book on database implementation might be helpful [347, 270]. Although alternative storage technologies exist and are being developed [653], databases are mostly stored on disks. Thus, we start out by introducing a simple disk model to capture I/O costs. Then, we say some words about database buffers, physical data organization, slotted pages and tuple identifiers (TIDs), physical record layout, physical algebra, and the iterator concept. These are the basic notions in order to start with the main purpose of this section: giving an overview over the possibilities available to structure the low level parts of a physical query evaluation plan. In order to calculate the I/O costs of these plan fragments, a more sophisticated cost model for several kinds of disk accesses is introduced.

4.1 Disk Drive

Figure 4.1 shows a top and a side view of a typical disk. A disk consists of several platters that rotate around the spindle at a fixed speed. The platters are coated with a magnetic material on at least one of their surfaces. All coated sides are organized into the same pattern of concentric circles. One concentric circle is called a track. All the tracks residing exactly underneath and above each other form a cylinder. We assume that there is only one read/write head for every coated surface.¹ All tracks of a cylinder can be accessed with only minor adjustments at the same time by their respective heads. By moving the arm around the arm pivot, other cylinders can be accessed. Each track is partitioned into sectors. Sectors have a disk specific (almost) fixed capacity of 512 B. The read and write granularity is a sector. Read and write accesses take place while the sector passes under the head.

The top view of Figure 4.1 shows that the outer sectors are longer than the

¹This assumption is valid for most but not all disks.




Figure 4.1: Disk drive assembly

inner sectors. The highest density (e.g. in bits per centimeter) at which bits can be separated is fixed for a given disk. For storing 512 B, this results in a minimum sector length which is used for the tracks of the innermost cylinder. Thus, since sectors on outer tracks are longer, storage capacity is wasted there. To overcome this problem, disks have a varying number of sectors per track. (This is where the picture lies.) Therefore, the cylinders are organized into zones. Every zone contains a fixed number of consecutive cylinders, each having a fixed number of sectors per track. Between zones, the number of sectors per track varies. Outer zones have more sectors per track than inner zones. Since the platters rotate with a fixed angular speed, sectors of outer cylinders can be read faster than sectors of inner cylinders. As a consequence, the throughput for reading and writing outer cylinders is higher than for inner cylinders.

Assume that we sequentially read all the sectors of all tracks of some consecutive cylinders. After reading all sectors of some track, we must proceed to the next track. If it is contained in the same cylinder, then we must (simply) use another head: a *head switch* occurs. Due to calibration, this takes some time. Thus, if all sectors start at the same angular position, we come too late to read the first sector of the next track and have to wait. To avoid this, the angular start positions of the sectors of tracks in the same cylinder are skewed such that this *track skew* compensates for the head switch time. If the next track is contained in another cylinder, the heads have to switch to the next cylinder. Again, this takes time and we miss the first sector if all sectors of a surface start at the same angular positions. *Cylinder skew* is used such that the time needed for this switch does not make us miss to start reading the next sector. In general, skewing works in only one direction.

A sector can be addressed by a triple containing its cylinder, head (surface), and sector number. This triple is called the physical address of a sector. However, disks are accessed using logical addresses. These are called *logical block numbers* (LBN) and are consecutive numbers starting with zero. The disk internally maps LBNs to physical addresses. This mapping is captured in the following table:

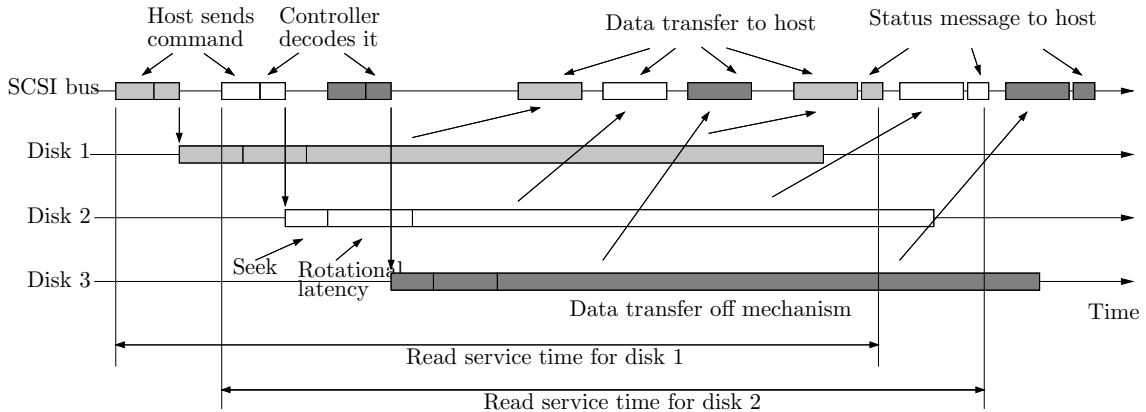


Figure 4.2: Disk drive read request processing

cylinder	track	LBN	number of sectors per track
0	0	0	573
	1	573	573
...
	5	2865	573
1	0	3438	573
...
15041	0	35841845	253
...

However, this ideal view is disturbed by the phenomenon of bad blocks. A *bad block* is one with a defect and it cannot be read or written. After a block with a certain LBN is detected to be bad, it is assigned to another sector. The above mapping changes. In order to be able to redirect LBNs, extra space on the disk must exist. Hence, some cylinders, tracks, and sectors are reserved for this purpose. They may be scattered all over the platters. Redirected blocks cause hiccups during sequential read.

Building (see e.g. [546]) and modeling (see e.g. [504, 640, 701, 702, 763, 803]) disk drives is challenging. Whereas the former is not really important when building query compiler, the latter is, as we have to attach costs to query evaluation plans. These costs reflect the amount of time we occupy the resource disk. Since disks are relatively slow, they may become the bottleneck of a database server. Modeling and minimizing disk access (time) is thus an important topic. Consider the case where we want to read a block from a SCSI disk. Simplified, the following actions take place and take their time (see also Fig. 4.2):

1. The host sends the SCSI command.
2. The disk controller decodes the command and calculates the physical address.
3. During the seek the disk drive's arm is positioned such that the accord-

ing head is correctly placed over the cylinder where the requested block resides. This step consists of several phases.

- (a) The disk controller accelerates the arm.
 - (b) For long seeks, the arm moves with maximum velocity (coast).
 - (c) The disk controller slows down the arm.
 - (d) The disk arm settles for the desired location. The settle times differ for read and write requests. For reads, an aggressive strategy is used. If, after all, it turns out that the block could not be read correctly, we can just discard it. For writing, a more conservative strategy is in order.
4. The disk has to wait until the sector where the requested block resides comes under the head (rotation latency).
 5. The disk reads the sector and transfers data to the host.
 6. Finally, it sends a status message.

Note that the transfers for different read requests are interleaved. This is possible since the capacity of the SCSI bus is higher than the read throughput of the disk. Also note that we did not mention the operating system delay and congestions on the SCSI bus.

Disk drives apply several strategies to accelerate the above-mentioned round-trip time and access patterns like sequential read. Among them are caching, read-ahead, and command queuing. (discuss interleaving?)

ToDo

The seek and rotation latency times highly depend on the head's position on the platter surface. Let us consider seek time. A good approximation of the seek time where d cylinders have to be travelled is given by

$$seektime(d) = \begin{cases} c_1 + c_2\sqrt{d} & d \leq c_0 \\ c_3 + c_4d & d > c_0 \end{cases}$$

where the constants c_i are disk-specific. The constant c_0 indicates the maximum number of cylinders where no coast takes place: seeking over a distance of more than c_0 cylinders results in a phase where the disk arm moves with maximum velocity.

For disk accesses, the database system must be able to estimate the time they take to be executed. First of all, we need the parameters of the disk. It is not too easy to get hold of them, but we can make use of several tools to extract them from a given disk [209, 266, 752, 660, 815, 816]. However, then we have a big problem: when calculating I/O costs, the query compiler has no idea where the head will be when the query evaluation plan emits a certain read (or write) command. Thus, we have to find another solution. In the following, we will discuss a rather simplistic cost model that will serve us to get a feeling for disk behavior. Later, we develop a more realistic model (Section 4.17).

The solution is rather trivial: we sum up all command sending and interpreting times as well the times for positioning (seek and rotation latency) which

form by far the major part. Let us call the result *latency time*. Then, we assume an average latency time. This, of course, may result in large errors for a single request. However, on average, the error can be as “low” as 35% [640]. The next parameter is the *sustained read rate*. The disk is assumed to be able to deliver a certain amount of bytes per second while reading data stored consecutively. Of course, considering multi-zone disks, we know that this is oversimplified, but we are still in our simplistic model. Analogously, we have a sustained write rate. For simplicity, we will assume that this is the same as the sustained read rate. Last, the capacity is of some interest. A hypothetical disk (inspired by disks available in 2004) then has the following parameters:

Model 2004		
Parameter	Value	Abbreviated Name
capacity	180 GB	D_{cap}
average latency time	5 ms	D_{lat}
sustained read rate	100 MB/s	D_{srr}
sustained write rate	100 MB/s	D_{swr}

The time a disk needs to read and transfer n bytes is then approximated by $D_{\text{lat}} + n/D_{\text{srr}}$. Again, this is overly simplistic: (1) due to head switches and cylinder switches, long reads have lower throughput than short reads and (2) multiple zones are not modelled correctly. However, let us use this very simplistic model to get some feeling for disk costs.

Database management system developers distinguish between *sequential* I/O and *random* I/O. For sequential I/O, there is only one positioning at the beginning and then, we can assume that data is read with the sustained read rate. For random I/O, one positioning for every unit of transfer—typically a page of say 8 KB—is assumed. Let us illustrate the effect of positioning by a small example. Assume that we want to read 100 MB of data stored consecutively on a disk. Sequential read takes 5 ms plus 1 s. If we read in blocks of 8 KB where each block requires positioning then reading 100 MB takes 65 s.

Assume that we have a relation of about 100 MB in size, stored on a disk, and we want to read it. Does it take 1 s or 65 s? If the blocks on which it is stored are randomly scattered on disk and we access them in a random order, 65 s is a good approximation. So let us assume that it is stored on consecutive blocks. Assume that we read in chunks of 8 KB. Then,

- other applications,
- other transactions, and
- other read operations of the same query evaluation plan

could move the head away from our reading position. (Congestion on the SCSI bus may also be problem.) Again, we could be left with 65 s. Reading the whole relation with one read request is a possibility but may pose problems to the buffer manager. Fortunately, we can read in chunks much smaller than 100 MB. Consider Figure 4.3. If we read in chunks of 100 8 KB blocks we are already pretty close to one second (within a factor of two).




Figure 4.3: Time to read 100 MB from disk (depending on the number of 8 KB blocks read at once)

Note that the interleaving of actions does not necessarily mean a negative impact. This depends on the point of view, i.e. what we want to optimize. If we want to optimize response time for a single query, then obviously the impact of concurrent actions is negative. If, however, we want to optimize resource (here: disk) usage, concurrent actions might help.

ToDo?

There are two important things to learn here. First, sequential read is much faster than random read. Second, the runtime system should secure sequential read. The latter point can be generalized: the runtime system of a database management system has, as far as query execution is concerned, two equally important tasks:

- allow for efficient query evaluation plans and
- allow for smooth, simple, and robust cost functions.

Typical measures on the database side are

- carefully chosen physical layout on disk
(e.g. cylinder or track-aligned extents [661, 662, 659], clustering),
- disk scheduling, multi-page requests
[195, 391, 669, 670, 677, 697, 727, 807, 814],
- (asynchronous) prefetching,
- piggy-back scans,
- buffering (e.g. multiple buffers, replacement strategy from [64] to [513]),
and last but not least
- efficient and robust algorithms for algebraic operators [294].

Let us take yet another look at it. 100 MB can be stored on 12800 8 KB pages. Figure 4.4 shows the time to read n random pages. In our simplistic cost model, reading 200 pages randomly costs about the same as reading 100 MB sequentially. That is, reading 1/64th of 100 MB randomly takes as long as reading the 100 MB sequentially. Let us denote by a the positioning time, s the sustained read rate, p the page size, and d some amount of consecutively stored bytes. Let us calculate the break-even point

$$\begin{aligned} n * (a + p/s) &= a + d/s \\ n &= (a + d/s)/(a + p/s) \\ &= (as + d)/(as + p) \end{aligned}$$

a and s are disk parameters and, hence, fixed. For a fixed d , the break-even point depends on the page size. This is illustrated in Figure 4.5. The x-axis is the page size p in multiples of 1 K and the y-axis is $(d/p)/n$ for $d = 100$ MB.

For sequential reads, the page size does not matter. (Be aware that our simplistic model heavily underestimates sequential reads.) For random reads,




Figure 4.4: Time needed to read n random pages




Figure 4.5: Break-even point in fraction of total pages depending on page size

as long as a single page is read, it matters neither: reading a single page of 1 KB lasts 5.0097656 ms, for an 8 KB page the number is 5.0781250 ms. From all this, we could draw the conclusion that the larger the page the better. However, this is only true for the disk, not, e.g., for the buffer or the SCSI bus. If we need to access only 500 B of a page, then the larger the page the higher the fraction that is wasted. This is not as severe as it sounds. Other queries or transactions might need other parts of the page during a single stay in the buffer. Let us call the fraction of the page that is read by some transaction during a stay in the buffer by utilization. Obviously, the higher the utilization the better is our usage of the main memory in which the buffer resides. For smaller pages, the utilization is typically higher than for larger pages. The frequency by which pages are used is another factor. [314, 315].

Excursion. Consider the root page of a B-tree. It is accessed quite frequently and most of its parts will be used, no matter how large it is. Hence, utilization is always good. Thus, the larger the root page of a B-tree the better. On the other hand, consider a leaf page of a B-tree that is much bigger than main memory. During a single stay of it, only a small fraction of the page will be used. That is, smaller leaf pages are typically better. By converting everything to money instead of time, Gray and Graefe [314] as well as Lomet [489] come to the conclusion that a page size between 8 and 16 KB was a good choice at the end of the last century.

For the less simplistic model of disk access costs developed in Section 4.17, we need to describe a disk drive by a set of parameters. These parameters are summarized in Table 4.1.

Let us close this section by giving upper bounds on seek time and rotational latency. Qyang proved the following theorem which gives a tight upper bound of disk seek time if several cylinders of a consecutive range of cylinders have to be visited [600].

Theorem 4.1.1 (Qyang) *If the disk arm has to travel over a region of C cylinders, it is positioned on the first of the C cylinders and has to stop at $s - 1$ of them, then $sD_{seek}(C/s)$ is an upper bound for the seek time.*

The time required for s consecutive sectors in a track of zone i to pass by the head is

$$D_{rot}(s, i) = sD_{Zscan}(i) = s \frac{D_{rot}}{D_{Zspt}(i)} \quad (4.1)$$

A trivial upper bound for the rotational delay is a full rotation.

4.2 Database Buffer

The database buffer

1. is a finite piece of memory,
2. typically supports a limited number of different page sizes (mostly one or two),

D_{cyl}	total number of cylinders
D_{track}	total number of tracks
D_{sector}	total number of sectors
D_{tpc}	number of tracks per cylinder (= number of surfaces)
D_{cmd}	command interpretation time
D_{rot}	time for a full rotation
D_{rdsettle}	time for settle for read
D_{wrsettle}	time for settle for write
D_{hdswitch}	time for head switch
D_{Zone}	total number of zones
$D_{\text{Zcyl}}(i)$	number of cylinders in zone i
$D_{\text{Zspt}}(i)$	number of sectors per track in zone i
$D_{\text{Zspc}}(i)$	number of sectors per cylinder in zone i ($= D_{\text{tpc}}D_{\text{Zspt}}(i)$)
$D_{\text{Zscan}}(i)$	time to scan a sector in zone i ($= D_{\text{rot}}/D_{\text{Zspt}}(i)$)
D_{avgseek}	average seek costs
D_{c_0}	parameter for seek cost function
D_{c_1}	parameter for seek cost function
D_{c_2}	parameter for seek cost function
D_{c_3}	parameter for seek cost function
D_{c_4}	parameter for seek cost function
$D_{\text{seek}}(d)$	cost of a seek of d cylinders
	$D_{\text{seek}}(d) = \begin{cases} D_{c_1} + D_{c_2}\sqrt{d} & \text{if } d \leq D_{c_0} \\ D_{c_3} + D_{c_4}d & \text{if } d > D_{c_0} \end{cases}$
$D_{\text{rot}}(s, i)$	rotation cost for s sectors of zone i ($= sD_{\text{Zscan}}(i)$)

Table 4.1: Disk drive parameters and elementary cost functions

3. is often fragmented into several buffer pools,
4. each having a replacement strategy (typically enhanced by hints).

Given the page identifier, the buffer frame is found by a hashtable lookup. Accesses to the hash table and the buffer frame need to be synchronized. Before accessing a page in the buffer, it must be fixed. These points account for the fact that the costs of accessing a page in the buffer are, therefore, greater than zero.

4.3 Physical Database Organization

We call everything that is stored in the database and relevant for answering queries a *database item*. Let us exclude meta data. In a relational system, a database item can be a relation, a fragment of a relation (if the relation is horizontally or vertically fragmented), a segment, an index, a materialized view,

or an index on a materialized view. In object-oriented databases, a database item can be the extent of a class, a named object, an index and so forth. In XML databases, a database item can be a named document, a collection of documents, or an index. Access operations to database items form the leaves of query evaluation plans.

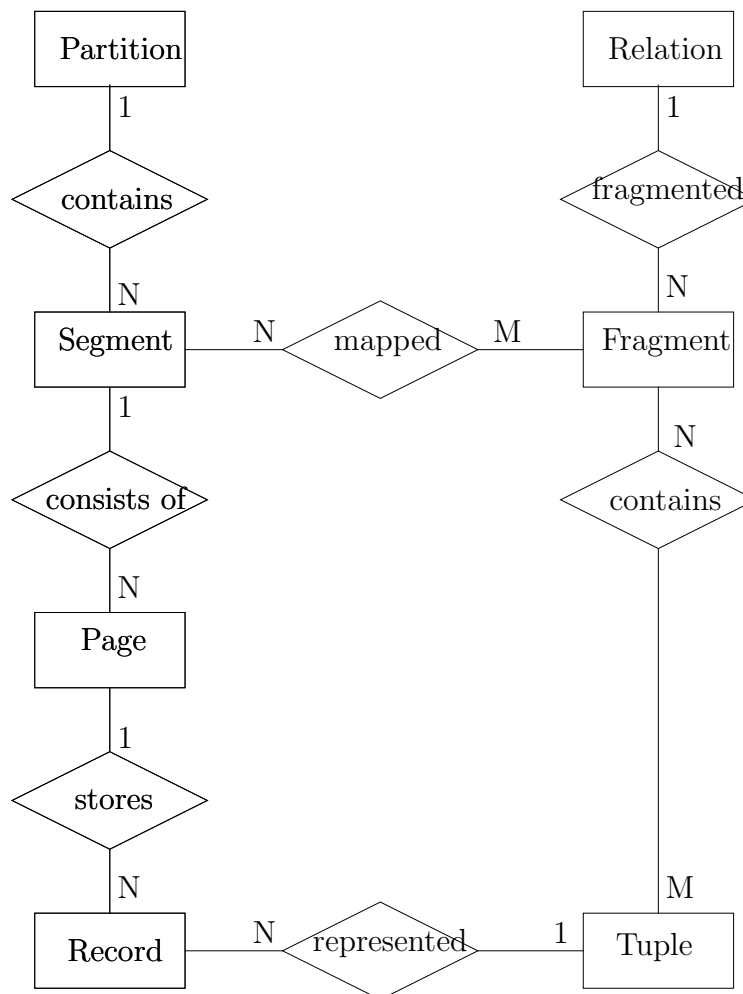


Figure 4.6: Physical organization of a relational database

The physical algebra implemented in the query execution engine of some runtime systems allow to access database items. Since most database items consist of several data items (tuples, objects, documents), these access operations produce a *stream* of data items. This kind of collection-valued access operation is called a *scan*. Consider the simple query

```

select *
from Student
  
```

This query is valid only if the database item (relation) **Student** exists. It could

be accessible via a `relation scan` operation `rscan(Student)`. However, in reality we have to consider the physical organization of the database.

Figure 4.6 gives an overview of how relations can be stored in a relational database system. Physical database items can be found on the left-hand side, logical database items on the right-hand side. A fraction of a physical disk is a partition. It can be an operating system file or a raw partition. A partition is organized into several segments. A segment consists of several pages. The pages within a segment are typically accessible by a non-negative integer in $[0, n[$, where n is the number of pages of the segment². Iterative access to all pages of a segment is typically possible. The access is called a *scan*. As there are several types of segments (e.g. data segments, index segments), several kinds of scans exist. Within a page, *physical records* are stored. Each physical record represents a (part of a) tuple of a fragment of a relation.

Fragments are mapped to segments and relations are partitioned into fragments. In the simplest and most common organization, every relation has only one fragment with a one-to-one mapping to segments, and for every tuple there exists exactly one record representing only this tuple. Hence, both of relationships *mapped* and *represented* are one-to-one. However, this organization does not scale well. A relation could be larger than a disk. Even if a large relation, say 180 GB fits on a disk, scanning it takes half an hour (Model 2004). Horizontal partitioning and allocation of the fragments on several disks reduces the scan time by allowing for parallelism. Vertical partitioning is another means of reducing I/O [176]. Here, a tuple is represented by several physical records, each one containing a subset of the tuple's attributes. Since the relationship *mapped* is N:M, tuples from different relations can be stored in the same segment. Furthermore, in distributed database systems some fragments might be stored redundantly at different locations to improve access times [112, 439, 601, 565]. Some systems support clustering of tuples of different relations. For example, department tuples can be clustered with employee tuples such that those employees belonging to the department are close together and close to their department tuple. Such an organization speeds up join processing.

To estimate costs, we need a model of a segment. We assume an extent-based implementation. That is, a segment consists of several extents³. Each extent occupies consecutive sectors on disk. For simplicity, we assume that whole cylinders belong to a segment. Then, we can model segments as follows. Each segment consists of a sequence of *extents*. Each extent is stored on *consecutive cylinders*. Cylinders are exclusively assigned to a segment. We then describe each extent j as a pair (F_j, L_j) where F_j is the first and L_j the last cylinder of a consecutive sequence of cylinders. A segment can then be described by a sequence of such pairs. We assume that these pairs are sorted in ascending order. In such a description, an extent may include a zone boundary. Since cost functions are dependent on the zone, we break up cylinder ranges that are not contained in a single zone. The result can be described by a sequence of triples

²This might not be true. Alternatively, the pages of a partition can be consecutively numbered.

³Extents are not shown in Fig. 4.6. They can be included between Partitions and Segments.




Figure 4.7: Slotted pages and TIDs

(F_i, L_i, z_i) where F_i and L_i mark a range of consecutive cylinders in a zone z_i . Although the z_i 's can be inferred from the cylinder numbers, we include them for clarity. Also of interest are the total number of sectors in a segment and the number of cylinders $S_{\text{cpe}}(i)$ in an extent i . Summarizing, we describe a segment by the parameter given in Table 4.2.

S_{ext}	number of extents in the segment
S_{sec}	total number of sectors in the segment $(= \sum_{i=1}^{S_{\text{ext}}} S_{\text{cpe}}(i) D_{\text{Zspc}}(S_{\text{zone}}(i)))$
$S_{\text{first}}(i)$	first cylinder in extent i
$S_{\text{last}}(i)$	last cylinder in extent i
$S_{\text{cpe}}(i)$	number of cylinders in extent i $(= S_{\text{last}}(i) - S_{\text{first}}(i) + 1)$
$S_{\text{zone}}(i)$	zone of extent i

Table 4.2: Segment parameters

4.4 Slotted Page and Tuple Identifier (TID)

Let us briefly review *slotted pages* and the concept of *tuple identifiers (TIDs)* (see Figure 4.7) [37, 36, 490, 734]. Sometimes, *record identifier* or *row identifier* (RID) is used in the literature. A TID consists of (at least) two parts. The first part identifies a page, the second part a slot on a *slotted page*. The slot contains—among other things, e.g. the record's size—a (relative) pointer to the actual record. This way, the record can be moved within the page without invalidating its TID. When a record grows beyond the available space, it is moved to another page and leaves a forward pointer (again consisting of a page and a slot identifier) in its original position. This happened to the TID [273, 1] in Figure 4.7. If the record has to be moved again, the forward pointer is adjusted. This way, at most two page accesses are needed to retrieve a record, given its TID. For evaluating the costs of record accesses, we will assume that




Figure 4.8: Various physical record layouts

the fraction of moved records is known.

4.5 Physical Record Layouts

A physical record represents a tuple, object, or some other logical entity or fraction thereof. In case it represents a tuple, it consists of several fields, each representing the value of an attribute. These values can be integers, floating point numbers, or strings. In case of object-oriented or object-relational systems, the values can also be of a complex type. Tuple identifiers are also possible as attribute values [622]. This can, for example, speed up join processing.

In any case, we can distinguish between types whose values all exhibit the same fixed length and those whose values may vary in length. In a physical record, the values of fixed-length attributes are concatenated and the offset from the beginning of the record to the value of some selected attribute can be inferred from the types of the values preceding it. This differs for values of varying length. Here, several encodings are possible. Some simple ones are depicted in Figure 4.8. The topmost record encodes varying length values as a sequence of pairs of the form $[size, value]$. This encoding has the disadvantage that access to an attribute of varying length is linear in the number of those preceding it. This disadvantage is avoided in the solution presented in the middle. Instead of storing the sizes of the individual values, there is an array containing relative offsets into the physical record. They point to the start of the values. The length of the values can be inferred from these offsets and, in case of the last value, from the total length of the physical record, which is typically stored in its slot. Access to a value of varying size is now simplified to an indirect memory access plus some length calculations. Although this might

be cheaper than the first solution, there is still a non-negligible cost associated with an attribute access.

The third physical record layout can be used to represent compressed attribute values and even compressed length information for parts of varying size. Note that if fixed size fields are compressed, their length becomes varying. Access to an attribute now means decompressing length/offset information and decompressing the value itself. The former is quite cheap: it boils down to an indirect memory access with some offset taken from a matrix [799]. The cost of the latter depends on the compression scheme used. It should be clear that accessing an attribute value now is even more expensive. To make the costs of an attribute access explicit was the sole purpose of this small section.

Remark Westmann et al. discuss an efficient implementation of compression and evaluate its performance [799]. Yiannis and Zobel report on experiments with several compression techniques used to speed up the sort operator. For some of them, the CPU usage is twice as large [?].

4.6 Physical Algebra (Iterator Concept)

Physical algebraic operators are mostly implemented as *iterators*. This means that they support the the interface operations `open`, `next`, and `close`. With `open`, the stream of items (e.g. tuples) is initialized. With `next`, the next item on the stream is fetched. When no more items are available, e.g. `next` returns false, `close` can be called to clean up things. The iterator concept is explained in many text books (e.g. [270, 347, 413]) and the query processing survey by Graefe [294]. This basic iterator concept has been extended to better cope with nested evaluation by Westmann in his thesis [797], Westmann et al. [799], and Graefe [298]. The two main issues are separation of storage allocation and initialization, and batched processing. The former splits `open` into resource allocation, initialization of the operator, and initialization of the iterator.

4.7 Simple Scan

Let us come back to the scan operations. A logical operation for scanning relations (which could be called `rscan`) is rarely supported by relational database management systems. Instead, they provide (physical) scans on segments. Since a (data) segment is sometimes called *file*, the correct plan for the above query is often denoted by `fscan(Student)`. Several assumptions must hold: the `Student` relation is not fragmented, it is stored in a single segment, the name of this segment is the same as the relation name, and no tuples from other relations are stored in this segment. Until otherwise stated, we will assume that relations are not partitioned, are stored in a single segment and that the segment can be inferred from the relation's name. Instead of `fscan(Student)`, we could then simply use `Student` to denote leaf nodes in a query execution plan. If we want to use a variable that is bound subsequently to each tuple in a relation, the query

```
select *
```

from Student

can be expressed as **Student**[s] instead of **Student**. In this notation, the output stream contains tuples having a single attribute s bound to a tuple. Physically, s will not hold the whole tuple but, for example, a pointer into the buffer where the tuple can be found. An alternative is a pointer to a slot of a slotted page contained in the buffer.

A simple scan is an example for a *building block*. In general, a building block is something that is used as a bottommost operator in a query evaluation plan. Hence, every leaf node of a query evaluation plan is a building block or a part thereof. This is not really a sharp definition, but is sometimes useful to describe the behavior of a query compiler: after their determination, it will leave building blocks untouched even if reorderings are hypothetically possible. Although a building block can be more than a leaf node (scan) of a query evaluation plan, it will never include more than a single database item. As soon as more database items are involved, we use the notion of *access path*, a term which will become more precise later on when we discuss index usage.

The disk access costs for a simple scan can be derived from the considerations in Section 4.1 and Section 4.17.

4.8 Scan and Attribute Access

Strictly speaking, a plan like $\sigma_{\text{age}>30}(\text{Student}[s])$ is invalid, since the tuple stream produced by **Student**[s] contains tuples with a single attribute s . We have a choice. Either we assume that attribute access takes place implicitly, or we make it explicit. Whether this makes sense or not depends on the database management system for which we generate plans. Let us discuss the advantages of explicit attribute retrieval. Assume **s.age** retrieves the age of a student. Then we can write $\sigma_{\text{s.age}>30}(\text{Student}[s])$, where there is some non-neglectable cost for **s.age**. The expression $\sigma_{\text{s.age}>30 \wedge \text{s.age}<40}(\text{Student}[s])$ executes **s.age** twice. This is a bad idea. Instead, we would like to retrieve it once and reuse it later.

This purpose is well-served by the *map* operator (χ). It adds new attributes to a given tuple and is defined as

$$\chi_{a_1:e_1, \dots, a_n:e_n}(e) := \{t \circ [a_1 : c_1, \dots, a_n : c_n] \mid t \in e, c_i = e_i(t) \forall (1 \leq i \leq n)\}$$

where \circ denotes tuple concatenation and the a_i must not be in $\mathcal{A}(e)$. (Remember that $\mathcal{A}(e)$ is the set of attributes produced by e .) Every input tuple t is extended by new attributes a_i , whose values are computed by evaluating the expression e_i , in which free variables (attributes) are bound to the attributes (variables) provided by t .

The above problem can now be solved by

$$\sigma_{\text{age}>30 \wedge \text{age}<40}(\chi_{\text{age:s.age}}(\text{Student}[s])).$$

In general, it is beneficial to load attributes as late as possible. The latest point at which all attributes must be read from the page is typically just before a pipeline breaker⁴.

To see why this is useful, consider the simple query

```
select name
from Student
where age > 30
```

The plan

$$\Pi_n(\chi_{n:s.name}(\sigma_{a>30}(\chi_{a:s.age}(\mathbf{Student}[s])))$$

makes use of this feature, while

$$\Pi_n(\sigma_{a>30}(\chi_{n:s.name,a:s.age}(\mathbf{Student}[s])))$$

does not. In the first plan the **name** attribute is only accessed for those students with age over 30. Hence, it should be cheaper to evaluate. If the database management system does not support this selective access mechanism, we often find the scan enhanced by a list of attributes that is projected and included in the resulting tuple stream.

In order to avoid copying attributes from their storage representation to some main memory representation, some database management systems apply another mechanism. They support the evaluation of some predicates directly on the storage representation. These are boolean expressions consisting of simple predicates of the form $A\theta c$ for attributes A , comparison operators θ , and constants c . Instead of a constant, c could also be the value of some attribute or expression thereof given that it can be evaluated before the access to A .

Predicates evaluable on the disk representation are called *SARGable* where *SARG* is an acronym for *search argument*. Note that SARGable predicates may also be good for index lookups. Then they are called *index SARGable*. In case they can not be evaluated by an index, they are called *data SARGable* [672, 750, 275].

Since relation or segment scans can evaluate predicates, we have to extend our notation for scans. Let I be a database item like a relation or segment. Then, $I[v;p]$ scans I , binds each item in I successively to v and returns only those items for which p holds. $I[v;p]$ is equivalent to $\sigma_p(I[v])$, but cheaper to evaluate. If p is a conjunction of predicates, the conjuncts should be ordered such that the attribute access cost reductions described above are reflected (for details see Chapter ??). Syntactically, we express this by separating the predicates by a comma as in `Student[s; age > 30, name like '%m%']`. If we want to make a distinction between SARGable and non-SARGable predicates, we write $I[v;p_s;p_r]$, with p_s being the SARGable predicate and p_r a non-SARGable predicate. Additional extensions like a projection list are also possible.

⁴The page on which the physical record resides must be fixed until all attributes are loaded. Hence, an earlier point in time might be preferable.

4.9 Temporal Relations

Scanning a temporal relation or segment also makes sense. Whenever the result of some (partial) query evaluation plan is used more than once, it might be worthwhile to materialize it in some temporary relation. For this purpose, a `tmp` operator evaluates its argument expression and stores the result relation in a temporary segment. Consider the following example query.

```

select  e.name, d.name
from    Emp e, Dept d
where   e.age > 30 and e.age < 40 and e.dno = d.dno

```

It can be evaluated by

$$\text{Dept}[d] \bowtie_{e.dno=d.dno}^{\text{nl}} \sigma_{e.age>30 \wedge e.age<40}(\text{Emp}[d]).$$

Since the inner (right) argument of the nested-loop join is evaluated several times (once for each department), materialization may pay off. The plan then looks like

$$\text{Dept}[d] \bowtie_{e.dno=d.dno}^{\text{nl}} \text{Tmp}(\sigma_{e.age>30 \wedge e.age<40}(\text{Emp}[d])).$$

If we choose to factorize and materialize a common subexpression, the query evaluation plan becomes a DAG. Alternatively, we could write a small “program” that has some statements materializing some expressions which are then used later on. The last expression in a program determines its result. For our example, the program looks as follows.

1. $R_{\text{tmp}} = \sigma_{e.age>30 \wedge e.age<40}(\text{Emp}[d]);$
2. $\text{Dept}[d] \bowtie_{e.dno=d.dno}^{\text{nl}} R_{\text{tmp}}[e]$

The disk costs of writing and reading temporary relations can be calculated using the considerations of Section 4.1.

4.10 Table Functions

A *table function* is a function that returns a relation [495]. An example is `Primes(int from, int to)`, which returns all primes between `from` and `to`, e.g. via a sieve-method. It can be used in any place where a relation name can occur. The query

```

select  *
from    TABLE(Primes(1,100)) as p

```

returns all primes between 1 and 100. The attribute names of the resulting relation are specified in the declaration of the table function. Let us assume that for `Primes` a single attribute `prime` is specified. Note that table functions may take parameters. This does not pose any problems, as long as we

know that `Primes` is a table function and we translate the above query into `Primes(1, 100)[p]`. Although this looks exactly like a table scan, the implementation and cost calculations are different.

Consider the following query where we extract the years in which we expect a special celebration of Anton's birthday.

```
select *
from Friends f,
     TABLE(Primes(
             CURRENT_YEAR, EXTRACT(YEAR FROM f.birthday) + 100)) as p
where f.name = 'Anton'
```

The result of the table function depends on our friend Anton. Hence, a join is no solution. Instead, we have to introduce a new kind of join, the *d-join* where the *d* stands for dependent. It is defined as

$$R < S > = \{t \circ s \mid t \in T, s \in S(t)\}.$$

The above query can now be evaluated as

$$\chi_{b:EXTRACT_YEAR(f.birthday)+100}(\sigma_{f.name='Anton'}(Friends[f])) < Primes(c, b)[p] >$$

where we assume that some global entity *c* holds the value of `CURRENT_YEAR`.

Let us do the above query for all friends. We just have to drop the `where` clause. Obviously, this results in many redundant computations of primes. At the SQL level, using the birthday of the youngest friend is beneficial:

```
select *
from Friends f,
     TABLE(Primes(
             CURRENT_YEAR, (select max(birthday) from Friends) + 100)) as p
where p.prime ≥ f.birthday
```

At the algebraic level, this kind of optimizations will be considered in Section ??.

Things can get even more involved if table functions can consume and produce relations, i.e. arguments and results can be relations.

ToDo?

Little can be said about the disk costs of table functions. They can be zero if the function is implemented such that it does not access any disks (files stored there), but it can also be very expensive if large files are scanned each time it is called. One possibility is to let the database administrator specify the numbers the query optimizer needs. However, since parameters are involved, this is not really an easy task. Another possibility is to measure the table function's behavior whenever it is executed, and learn about its resource consumption.

4.11 Indexes

There exists a plethora of different index structures. In the context of relational database management systems, the most versatile and robust index is the B-tree or variants/improvements thereof (e.g. []). It is implemented in almost




Figure 4.9: Clustered vs. non-clustered index

every commercial database management system. Some support hash-indexes (e.g. []). Other data models or specific applications need specialized indexes. There exist special index structures for indexing path expressions in object-oriented databases (e.g. []) and XML databases (e.g. []). Special purpose indexes include join indexes (e.g. [345, 772]) multi-dimensional indexes (e.g. []), variant (projection) indexes [551], small materialized aggregates [527], bitmap indexes [], and temporal indexes (e.g. []). We cannot discuss all indexes and their exploitations for efficient query evaluation. This fills more than a single book. Instead, we concentrate on B-tree indexes. In general, a B-tree can be used to index several relations. We only discuss cases where B-trees index a single relation.

The *search key* (or *key* for short) of an index is the sequence of attributes of the *indexed relation* over which the index is defined. A key is a *simple key* if it consists of a single attribute. Otherwise, it is a *complex key*. Each entry in the B-tree's leaf page consists of pairs containing the key values and a sequence of tuple identifiers (typically sorted by increasing page number). Every tuple with a TID in this list satisfies the condition that its indexed attribute's values are equal to the key values. If for every sequence of key values there is at most one such tuple, we have a *unique* index, otherwise a *non-unique* index.

The leaf entries may contain values from additional (non-key) attributes. Then we call the index *attribute data added* and the additional attributes *data attributes*. If the index contains all attributes of the indexed relation—in its key or data attributes—storing the relation is no longer necessary. The result is an *index-only relation*. In this case, the concept of tuple identifiers is normally no longer used since tuples can now be moved frequently, e.g. due to a leaf page split. This has two consequences. First, the data part does not longer contain the TID. Second, other indexes on the index-only relation cannot have tuple identifiers in their data part either. Instead, they use the key of the index-only relation to uniquely reference a tuple. For this to work, we must have a unique index.

B-trees can be either *clustered* or *non-clustered* indexes. In a clustered index, the tuple identifiers in the list of leaf pages are ordered according to their page

numbers. Otherwise, it is a *non-clustered* index⁵. Figure 4.9 illustrates this. Range queries result in sequential access for clustered indexes and in random access for non-clustered indexes.

4.12 Single Index Access Path

4.12.1 Simple Key, No Data Attributes

Consider the *exact match query*

```
select name
from Emp
where eno = 1077
```

If there exists a unique index on the key attribute `eno`, we can first access the index to retrieve the TID of the employee tuple satisfying `eno = 1077`. Another page access yields the tuple itself which constitutes the result of the query. Let Emp_{eno} be the index on `eno`, then we can descend the B-tree, using 1077 as the search key. A predicate that can be used to descend the B-tree or, in general, governing search within an index structure, is called an *index sargable predicate*.

For the example query, the index scan, denoted as $\text{Emp}_{\text{eno}}[x; \text{eno} = 1077]$, retrieves a single leaf node entry with attributes `eno` and TID. Similar to the regular scan, we assume x to be a variable holding a pointer to this index entry. We use the notations $x.\text{eno}$ and $x.\text{TID}$ to access these attributes. To dereference the TID, we use the map (χ) operator and a dereference function `deref` (or `*` for short). It turns a TID into a pointer in the buffer area. This of course requires the page to be loaded, if it is not in the buffer yet. The complete plan for the query is

$$\Pi_{\text{name}}(\chi_{e:*(x.\text{TID}),\text{name}:e.\text{name}}(\text{Emp}_{\text{eno}}[x; \text{eno} = 1077]))$$

where we computed several new attributes with one χ operator. Note that they are dependent on previously computed attributes and, hence, the order of evaluation does matter.

We can make the dependency of the map operator more explicit by applying a d-join. Denote by \square an operator that returns a single empty tuple. Then

$$\Pi_{\text{name}}(\text{Emp}_{\text{eno}}[x; \text{eno} = 1077] \lt \chi_{e:*(x.\text{TID}),\text{name}:e.\text{name}}(\square) \gt)$$

is equivalent to the former plan. Joins and indexes will be discussed in Section 4.14.

A *range query* like

```
select name
from Emp
where age ≥ 25 and age ≤ 35
```

⁵Of course, any degree of clusteredness may occur and has to be taken into account in cost calculations.

specifies a range for the indexed attribute. It is evaluated by an index scan with *start* and *stop* conditions. In our case, the start condition is `age ≥ 25`, and the stop condition is `age ≤ 35`. The start condition is used to retrieve the first tuple satisfying it by searching within the B-tree. In our case, 25 is used to descend from the root to the leaf page containing the key 25. Then, all records with keys larger than 25 within the page are searched. Since entries in B-tree pages are sorted on key values, this is very efficient. If we are done with the leaf page that contains 25 and the stop key has not been found yet, we proceed to the next leaf page. This is possible since leaf pages of B-trees tend to be chained. Then all records of the next leaf page are scanned and so on until we find the stop key. The complete plan then is

$$\Pi_{\text{name}}(\chi_{e:*(x.\text{TID}),\text{name}:e.\text{name}}(\text{Emp}_{\text{age}}[x; 25 \leq \text{age}; \text{age} \leq 35]))$$

If the index on `age` is non-clustered, this plan results in random I/O. We can turn random I/O into sequential I/O by sorting the result of the index scan on its TID attribute before dereferencing it⁶. This results in the following plan:

$$\Pi_{\text{name}}(\chi_{e:*(\text{TID}),\text{name}:e.\text{name}}(\text{Sort}_{\text{TID}}(\text{Emp}_{\text{age}}[x; 25 \leq \text{age}; \text{age} \leq 35; \text{TID}])))$$

Here, we explicitly included the TID attribute of the index into the projection list.

Consider a similar query which demands the output to be sorted:

```

select    name, age
from      Emp
where     age ≥ 25 and age ≤ 35
order by age

```

Since an index scan on a B-tree outputs its result ordered on the indexed attribute, the following plan produces the perfect result:

$$\Pi_{\text{name,age}}(\chi_{e:*(x.\text{TID}),\text{name}:e.\text{name}}(\text{Emp}_{\text{age}}[x; 25 \leq \text{age}; \text{age} \leq 35]))$$

On a clustered index this is most probably the best plan. On a non-clustered index, random I/O disturbs the picture. We avoid that by sorting the result of the index scan on the TID attribute and, after accessing the tuples, restore the order on `age` as in the following plan:

$$\Pi_{\text{name,age}}(\text{Sort}_{\text{age}}(\chi_{e:*(\text{TID}),\text{name}:e.\text{name}}(\text{Sort}_{\text{TID}}(\text{Emp}_{\text{age}}[x; 25 \leq \text{age}; \text{age} \leq 35; \text{TID}]))))$$

An alternative to this plan is not to sort on the original indexed attribute (`age` in our example), but to introduce a new attribute that holds the rank in the sequence derived from the index scan. This leads to the plan

$$\Pi_{\text{name,age}}(\text{Sort}_{\text{rank}}(\chi_{e:*(\text{TID}),\text{name}:e.\text{name}}(\text{Sort}_{\text{TID}}(\chi_{\text{rank}:counter++}(\text{Emp}_{\text{age}}[x; 25 \leq \text{age}; \text{age} \leq 35; \text{TID}]))))))$$

⁶This might not be necessary, if `Emp` fits main memory. Then, preferably asynchronous I/O should be used.

This alternative might turn out to be more efficient since sorting on an attribute with a dense domain can be implemented efficiently. (We admit that in the above example this is not worth considering.) There is another important application of this technique: XQuery often demands output in document order. If this order is destroyed during processing, it must at the latest be restored when the output it produced [507]. Depending on the implementation (i.e. the representation of document nodes or their identifiers), this might turn out to be a very expensive operation.

The fact that index scans on B-trees return their result ordered on the indexed attributes is also very useful if a merge-join on the same attributes (or a prefix thereof, see Chapter 23 for further details) occurs. An example follows later on.

Some *predicates* are not index SARGable, but can still be evaluated with the index as in the following query

```
select name
from Emp
where age ≥ 25 and age ≤ 35 and age ≠ 30
```

The predicate `age ≠ 30` is an example of a *residual predicate*. We can once more extend the index scan and compile the query into

$$\Pi_{\text{name}}(\chi_{t:x.\text{TID},e:*t,\text{name}:e.\text{name}}(\text{Emp}_{\text{age}}[x; 25 \leq \text{age}; \text{age} \leq 35; \text{age} \neq 30]))$$

Some index scan implementations allow exclusive bounds for start and stop conditions. With them, the query

```
select name
from Emp
where age > 25 and age < 35
```

can be evaluated using

$$\Pi_{\text{name}}(\chi_{t:x.\text{TID},e:*t,\text{name}:e.\text{name}}(\text{Emp}_{\text{age}}[x; 25 < \text{age}; \text{age} < 35]))$$

If this is not the case, two residual predicates must be used as in

$$\Pi_{\text{name}}(\chi_{t:x.\text{TID},e:*t,\text{name}:e.\text{name}}(\text{Emp}_{\text{age}}[x; 25 \leq \text{age}; \text{age} \leq 35; \text{age} \neq 25, \text{age} \neq 35]))$$

Especially for predicates on strings, this might be expensive.

Start and stop conditions are optional. To evaluate

```
select name
from Emp
where age ≥ 60
```

we use `age ≥ 60` as the start condition to find the leaf page containing the key 60. From there on, we scan all the leaf pages “to the right”.

If we have no start condition, as in

```

select name
from Emp
where age  $\leq$  20

```

we descend the B-tree to the “leftmost” page, i.e. the page containing the smallest key value, and then proceed scanning leaf pages until we encounter the key 20.

Having neither a start nor stop condition is also quite useful. The query

```

select count(*)
from Emp

```

can be evaluated by counting the entries in the leaf pages of a B-tree. Since a B-tree typically occupies far fewer pages than the original relation, we have a viable alternative to a relation scan. The same applies to the aggregate functions `sum` and `avg`. The other aggregate functions `min` and `max` can be evaluated much more efficiently by descending to the leftmost or rightmost leaf page of a B-tree. This can be used to answer queries like

```

select min/max(salary)
from Emp

```

much more efficiently than by a relation scan. Consider the query

```

select name
from Emp
where salary = (select max(salary)
                 from Emp)

```

It can be evaluated by first computing the maximum salary and then retrieving the employees earning this salary. This requires two descendants into the B-tree, while obviously one is sufficient. Depending on the implementation of the index (scan), we might be able to perform this optimization.

Further, the result of an index scan, whether it uses start and/or stop conditions or not, is always sorted on the key. This property can be useful for queries with no predicates. If we have neither a start nor a stop condition, the resulting scan is called *full index scan*. As an example consider the query

```

select salary
from Emp
order by salary

```

which is perfectly answered by the following full index scan:

$\mathbf{Emp}_{\text{salary}}$

So far, we have only seen indexes on numerical attributes.

```

select name, salary
from Emp
where name  $\geq$  'Maaa'

```

gives rise to a start condition ' $\text{Maaa} \leq \text{name}$ '. From the query

```

select name, salary
from Emp
where name like 'M%'

```

we can deduce the start condition ' $\text{M} \leq \text{name}$ '.

To express all the different alternatives of index usage, we need a powerful (and runtime system dependent) index scan expression. Let us first summarize what we can specify for an index scan:

1. the name of the variable for index entries (or pointers to them),
2. the start condition,
3. the stop condition,
4. a residual predicate, and
5. a projection list.

A projection list has entries of the form $a : x.b$ for attribute names a and b and x being the name of the variable for the index entry. $a : x.a$ is also allowed and often abbreviated as a . We also often summarize start and stop conditions into a single expression like in $25 \leq \text{age} \leq 35$.

For a full index specification, we list all items in the subscript of the index name separated by a semicolon. Still, we need some extensions to express the queries with aggregation. Let a and b be attribute names, then we allow entries of the form $b : \text{aggr}(a)$ in the projection list and start/stop conditions of the form $\text{min}/\text{max}(a)$. The latter tells us to minimize/maximize the value of the indexed attribute a . Only a complete enumeration gives us the full details. On the other hand, extracting start and stop conditions and residual predicates from a given boolean expression is rather simple. Hence, we often summarize these three under a single predicate. This is especially useful when talking about index scans in general. If we have a full index scan, we leave out the predicate. We use a star '*' as an abbreviated projection list that projects all attributes of the index. (So far, these are the key attribute and the TID.) If the projection list is empty, we assume that only the variable/attribute holding the pointer to the index entry is projected.

Using this notation, we can express some plan fragments. These fragments are complete plans for the above queries, except that the final projection is not present. As an example, consider the following fragment:

$$\chi_{e:*TID, \text{name}:e.\text{name}}(\text{Emp}_{\text{salary}}[x; \text{TID}, \text{salary}])$$

All the plan fragments seen so far are examples of access paths. An *access path* is a plan fragment with building blocks concerning a single database item.

Hence, every building block is an access path. The above plans touch two database items: a relation and an index on some attribute of that relation. If we say that an index concerns the relation it indexes, such a fragment is an access path. For relational systems, the most general case of an access path uses several indexes to retrieve the tuples of a single relation. We will see examples of these more complex access paths in the following section. An access to the original relation is not always necessary. A query that can be answered solely by accessing indexes is called an *index only query*.

A query with `in` like

```
select  name
from    Emp
where  age in {28, 29, 31, 32}
```

can be evaluated by taking the minimum and the maximum found in the left-hand side of `in` as the start and stop conditions. We further need to construct a residual predicate. The residual predicate can be represented either as $\text{age} = 28 \vee \text{age} = 29 \vee \text{age} = 31 \vee \text{age} = 32$ or as $\text{age} \neq 30$.

An alternative is to use a d-join. Consider the example query

```
select  name
from    Emp
where  salary in {1111, 11111, 111111}
```

Here, the numbers are far apart and separate index accesses might make sense. Therefore, let us create a temporary relation `Sal` equal to $\{[s : 1111], [s : 11111], [s : 111111]\}$. When using it, the access path becomes

$$\text{Sal}[S] < \chi_{e:*TID, \text{name}:e.\text{name}}(\text{Emp}_{\text{salary}}[x; \text{salary} = S.s; TID]) >$$

Some B-tree implementations allow efficient searches for multiple ranges and implement *gap skipping* [31, 32, 144, 275, 276, 407, 467]. *Gap skipping*, sometimes also called *zig-zag skipping*, continues the search for keys in a new key range from the latest position visited. The implementation details vary but the main idea of it is that after one range has been completely scanned, the current (leaf) page is checked for its highest key. If it is not smaller than the lower bound of the next range, the search continues in the current page. If it is smaller than the lower bound of the next range, alternative implementations are described in the literature. The simplest is to start a new search from the root for the lower bound. Another alternative uses parent pointers to go up a page as long as the highest key of the current page is smaller than the lower bound of the next range. If this is no longer the case, the search continues downwards again.

Gap skipping gives even more opportunities for index scans and allows efficient implementations of various index nested loop join strategies.

4.12.2 Complex Keys and Data Attributes

In general, an index can have a complex key comprised of the key attributes k_1, \dots, k_n and the data attributes d_1, \dots, d_m . One possibility is to use a full

index scan on such an index. Having more attributes in the index makes it more probable that queries are index-only.

Besides a full index scan, the index can be descended to directly search for the desired tuple(s). Let us take a closer look at this possibility.

If the search predicate is of the form

$$k_1 = c_1 \wedge k_2 = c_2 \wedge \dots \wedge k_j = c_j$$

for some constants c_i and some $j \leq n$, we can generate the start and stop condition

$$k_1 = c_1 \wedge \dots \wedge k_j = c_j.$$

This simple approach is only possible if the search predicates define values for all search key attributes, starting from the first search key and then for all keys up to the j -th search key with no key attribute unspecified in between. Predicates concerning the other key attributes after the first non-specified key attribute and the additional data attributes only allow for residual predicates. This condition is often not necessary for multi-dimensional index structures, whose discussion is beyond the book.

With ranges things become more complex and highly dependent on the implementation of the facilities of the B-tree. Consider a query predicate restricting key values as follows

$$k_1 = c_1 \wedge k_2 \geq c_2 \wedge k_3 = c_3$$

Obviously, we can generate the start condition $k_1 = c_1 \wedge k_2 \geq c_2$ and the stop condition $k_1 = c_1$. Here, we neglected the condition on k_3 which becomes a residual predicate. However, with some care we can extend the start condition to $k_1 = c_1 \wedge k_2 \geq c_2 \wedge k_3 = c_3$: we only have to keep $k_3 = c_3$ as a residual predicate, since for k_2 values larger than c_2 , values different from c_3 can occur for k_3 .

If closed ranges are specified for a prefix of the key attributes as in

$$a_1 \leq k_1 \leq b_1 \wedge \dots \wedge a_j \leq k_j \leq b_j$$

we can generate the start key $k_1 = a_1 \wedge \dots \wedge k_j = a_j$, the stop key $k_1 = b_1 \wedge \dots \wedge k_j = b_j$, and

$$a_2 \leq k_2 \leq b_2 \wedge \dots \wedge a_j \leq k_j \leq b_j$$

as the residual predicate. If for some search key attribute k_j the lower bound a_j is not specified, the start condition cannot contain k_j and any k_{j+i} . If for some search key attribute k_j the upper bound b_j is not specified, the stop condition cannot contain k_j and any k_{j+i} .

Two further enhancements of the B-tree functionality possibly allow for alternative start/stop conditions:

- The B-tree implementation allows to specify the order (ascending or descending) for each key attribute individually.

- The B-tree implementation implements forward and backward scans (e.g. implemented in Rdb [31]).

So far, we are only able to exploit query predicates which specify value ranges for a prefix of all key attributes. Consider querying a person on his/her height and his/her hair color: `haircolor = 'blond' and height between 180 and 190`. If we have an index on `sex, haircolor, height`, this index cannot be used by means of the techniques described so far. However, since there are only the two values `male` and `female` available for `sex`, we can rewrite the query predicate to `(sex = 'm' and haircolor = 'blond' and height between 180 and 190) or (sex = 'f' and haircolor = 'blond' and height between 180 and 190)` and use two accesses to the index. This approach works fine for attributes with a small domain and is described by Antoshenkov [32]. (See also the previous section for gap skipping.) Since the possible values for key attributes may not be known to the query optimizer, Antoshenkov goes one step further and shifts the construction of search ranges to index scan time. Therefore, the index can be provided with a complex boolean expression which is then refined (rewritten) as soon as search key values become known. Search ranges are then generated dynamically, and gap skipping is applied to skip the intervals between the qualifying ranges during the index scan.

4.13 Multi Index Access Path

We wish to buy a used digital camera and state the following query:

```
select *
from Camera
where megapixel > 5 and distortion < 0.05
      and noise < 0.01
      and zoomMin < 35 and zoomMax > 105
```

We assume that on every attribute used in the `where` clause there exists an index. Since the predicates are conjunctively connected, we can use a technique called *index and-ing*. Every index scan returns a set (list) of tuple identifiers. These sets/lists are then intersected. This operation is also called *And merge* [485]. Using index and-ing, a possible plan is

```
((((
  Camera_megapixel[c; megapixel > 5; TID]
  ∩
  Camera_distortion[c; distortion < 0.05; TID])
  ∩
  Camera_noise[c; noise < 0.01; TID])
  ∩
  Camera_zoomMin[c; zoomMin < 35; TID])
  ∩
  Camera_zoomMax[c; zoomMax > 105; TID])
```

This results in a set of tuple identifiers that only needs to be dereferenced to access the according **Camera** tuples and produce the final result.

Since the costs of the expression clearly depend on the costs of the index scans and the size of the intermediate TID sets, two questions arise:

- In which order do we intersect the TID sets resulting from the index scans?
- Do we really apply all indexes before dereferencing the tuple identifiers?

The answer to the latter question is clearly “no”, if the next index scan is more expensive than accessing the records in the current TID list. It can be shown that the indexes in the cascade of intersections are ordered on increasing $(f_i - 1)/c_i$ terms, where f_i is the selectivity of the index and c_i its access cost. Further, we can stop as soon as accessing the original tuples in the base relation becomes cheaper than intersecting with another index and subsequently accessing the base relation.

EX

Index or-ing is used to process disjunctive predicates. Here, we take the *union* of the TID sets to produce a set of TIDs containing references to all qualifying tuples. Note that duplicates must be eliminated during the processing of the union. This operation is also called *Or merge* [485]. Consider the query

```
select *
from Emp
where yearsOfEmployment ≥ 30
      or age ≥ 65
```

This query can be answered by constructing a TID set using the expression

$$\text{Emp}_{\text{yearsOfEmployment}}[c; \text{yearsOfEmployment} \geq 30; \text{TID}] \cup \text{Emp}_{\text{age}}[c; \text{age} \geq 65; \text{TID}]$$

and then dereferencing the list of tuple identifiers. Again, the index accessing can be ordered for better performance. Given a general boolean expression in **and** and **or**, constructing the optimal access path using index and-ing and or-ing is a challenging task that will be discussed in Chapter ???. This task is even more challenging, if some simple predicates occur more than once in the complex boolean expression and factorization has to be taken into account. This issue was first discussed by Chaudhuri, Ganesan and Saragawi [123]. We will come back to this in Chapter ??.

The names *index and-ing* and *or-ing* become clear if bitmap indexes are considered. Then the bitwise **and** and **or** operations can be used to efficiently compute the intersection and union.

ToDo

Excursion on bitmap indexes. □

There are even more possibilities to work with TID sets. Consider the query

```
select *
from Emp
where yearsOfEmployment ≠ 10
      and age ≥ 65
```

This query can be evaluated by scanning the index on `age` and then eliminating all employees with `yearsOfEmployment = 10`:

$\text{Emp}_{\text{age}}[c; \text{age} \geq 65; \text{TID}] \setminus \text{Emp}_{\text{yearsOfEmployment}}[c; \text{yearsOfEmployment} \neq 10; \text{TID}]$

Let us call the application of set difference on index scan results *index differencing*.

Some predicates might not be very restrictive in the sense that more than half the index has to be scanned. By negating these predicates and using index differencing, we can make sure that at most half of the index needs to be scanned. As an example consider the query

```
select *
from Emp
where yearsOfEmployment ≤ 5
      and age ≤ 65
```

Assume that most of our employees' age is below 65. Then

$\text{Emp}_{\text{yearsOfEmployment}}[c; \text{yearsOfEmployment} \leq 5; \text{TID}] \setminus \text{Emp}_{\text{age}}[c; \text{age} > 65; \text{TID}]$

could be more efficient than

$\text{Emp}_{\text{yearsOfEmployment}}[c; \text{yearsOfEmployment} \leq 5; \text{TID}] \cap \text{Emp}_{\text{age}}[c; \text{age} \leq 65; \text{TID}]$

4.14 Indexes and Joins

There are two issues when discussing indexes and joins. The first is that indexes can be used to speed up join processing. The second is that index accesses can be expressed as joins. We discuss both of these issues, starting with the latter.

In our examples, we used the map operation to (implicitly) access the relation by dereferencing the tuple identifiers. We can make the implicit access explicit by exchanging the map operator by a d-join or even a join. Then, for example,

$\chi_{e:*TID, \text{name}:e.\text{name}}(\text{Emp}_{\text{salary}}[x; 25 \leq \text{age} \leq 35; \text{TID}])$

becomes

$\text{Emp}_{\text{salary}}[x; 25 \leq \text{age} \leq 35; \text{TID}] \lt \chi_{e:*TID, \text{name}:e.\text{name}}(\square) \gt$

where \square returns a single empty tuple. Assume that every tuple contains an attribute `TID` containing its TID. This attribute does not have to be stored explicitly but can be derived. Then, we have the following alternative access path for the join (ignoring projections):

$\text{Emp}_{\text{salary}}[x; 25 \leq \text{age} \leq 35] \bowtie_{x.\text{TID}=e.\text{TID}} \text{Emp}[e]$

ToDo

For the join operator, the *pointer-based join* implementation developed in the context of object-oriented databases may be the most efficient way to evaluate the access path [695]. Obviously, sorting the result of the index scan on the tuple identifiers can speed up processing since it turns random into sequential I/O. However, this destroys the order on the key which might itself be useful later on during query processing or required by the query⁷. Sorting the tuple identifiers was proposed by, e.g., Yao [836], Makinouchi, Tezuka, Kitakami, and Adachi in the context of RDB/V1 [496]. The different variants (whether or not and where to sort, join order) can now be transparently determined by the plan generator: no special treatment is necessary. Further, the join predicates can not only be on the tuple identifiers but also on key attributes. This often allows to join with other than the indexed relations (or their indexes) before accessing the relation.

Rosenthal and Reiner proposed to use joins to represent access paths with indexes [629]. This approach is very elegant since no special treatment for index processing is required. However, if there are many relations and indexes, the search space might become very large, as every index increases the number of joins to be performed. This is why Mohan, Haderle, Wang, and Cheng abandoned this approach and sketched a heuristics which determines an access path in case multiple indexes on a single table exist [532].

The query

```

select  name,age
from    Person
where   name like 'R%' and age between 40 and 50

```

is an index only query (assuming indexes on **name** and **age**) and can be translated to

$$\Pi_{\text{name,age}}(\text{Emp}_{\text{age}}[a; 40 \leq \text{age} \leq 50; \text{TIDa, age}] \bowtie_{\text{TIDa}=\text{TIDn}} \text{Emp}_{\text{name}}[n; \text{name} \geq' R'; \text{name} \leq' R'; \text{TIDn, name}])$$

Let us now discuss the former of the two issues mentioned in the section's introduction. The query

```

select  *
from    Emp e, Dept d
where   e.name = 'Maier' and e.dno = d.dno

```

can be directly translated to

$$\sigma_{e.\text{name}='Maier'}(\text{Emp}[e]) \bowtie_{e.\text{dno}=d.\text{dno}} \text{Dept}[d]$$

⁷Restoring the order may be cheaper than typical sorting since tuples can be numbered before the first sort on tuple identifiers, and this dense numbering leads to efficient sort algorithms.

If there are indexes on `Emp.name` and `Dept.dno`, we can replace $\sigma_{e.name='Maier'}(\text{Emp}[e])$ by an index scan as we have seen previously:

$$\chi_{e:*(x.TID),\mathcal{A}(\text{Emp}):e.*}(\text{Emp}_{\text{name}}[x; \text{name} = \text{'Maier'}])$$

Here, $\mathcal{A}(\text{Emp}) : t.*$ abbreviates access to all `Emp` attributes. This especially includes `dno:t.dno`. (Strictly speaking, we do not have to access the `name` attribute, since its value is already known.)

As we have also seen, an alternative is to use a d-join instead:

$$\text{Emp}_{\text{name}}[x; \text{name} = \text{'Maier'}] < \chi_{t:*(x.TID),\mathcal{A}(e)t.*}(\square) >$$

Let us abbreviate $\text{Emp}_{\text{name}}[x; \text{name} = \text{'Maier'}]$ by E_i and $\chi_{t:*(x.TID),\mathcal{A}(e)t.*}(\square)$ by E_a .

Now, for any `e.dno`, we can use the index on `Dept.dno` to access the according department tuple:

$$E_i < E_a > < \text{Dept}_{\text{dno}}[y; y.dno = dno] >$$

Note that the inner expression $\text{Dept}_{\text{dno}}[y; y.dno = dno]$ contains the free variable `dno`, which is bound by E_a . Dereferencing the TID of the department results in the following algebraic modelling which models a complete *index nested loop join*:

$$E_i < E_a > < \text{Dept}_{\text{dno}}[y; y.dno = \text{dno}; \text{dTID} : y.TID] > < \chi_{u:*\text{dTID},\mathcal{A}(\text{Dept})u.*}(\square) >$$

Let us abbreviate $\text{Dept}_{\text{dno}}[y; y.dno = \text{dno}; \text{dTID} : y.TID]$ by D_i and $\chi_{u:*\text{dTID},\mathcal{A}(\text{Dept})u.*}(\square)$ by D_a . Fully abbreviated, the expression then becomes

$$E_i < E_a > < D_i > < D_a >$$

Several optimizations can possibly be applied to this expression. Sorting the *outer* of a d-join is useful under several circumstances since it may

- turn random I/O into sequential I/O and/or
- avoid reading the same page twice.

In our example expression,

- we can sort the result of expression E_i on TID in order to turn random I/O into sequential I/O, if there are many employees named “Maier”.
- we can sort the result of the expression $E_i < E_a >$ on `dno` for two reasons:
 - If there are duplicates for `dno`, i.e. there are many employees named “Maier” in each department, then this guarantees that no index page (of the index `Dept.dno`) has to be read more than once.
 - If additionally `Dept.dno` is a clustered index or `Dept` is an index-only table contained in `Dept.dno`, then large parts of the random I/O can be turned into sequential I/O.

- If the result of the inner is materialized (see below), then only one result needs to be stored. Note that sorting is not necessary, but grouping would suffice to avoid duplicate work.
- We can sort the result of the expression $E_i < E_a > < D_i >$ on dTID for the same reasons as mentioned above for sorting the result of E_i on TID.

EX

The reader is advised to explicitly write down the alternatives. Another exercise is to give plan alternatives for the different cases of DB2's Hybrid Join [275] which can now be decomposed into primitives like relation scan, index scan, d-join, sorting, TID dereferencing, and access to a unique index (see below).

Let us take a closer look at materializing the result of the inner of the d-join. IBM's DB2 for MVS considers temping (i.e. creating a temporary relation) the inner if it is an index access [275]. Graefe provides a general discussion on the subject [298]. Let us start with the above example. Typically, many employees will work in a single department and possibly several of them are called "Maier". For everyone of them, we can be sure that there exists at most one department. Let us assume that referential integrity has been specified. Then, there exists exactly one department for every employee. We have to find a way to rewrite the expression

$$E_i < E_a > < \text{Dept}_{\text{dno}}[y; y.\text{dno} = \text{dno}; \text{dTID} : y.\text{TID}] >$$

such that the mapping $\text{dno} \longrightarrow \text{dTID}$ is explicitly materialized (or, as one could also say, *cached*). For this purpose, Hellerstein and Naughton introduced a modified version of the map operator that materializes its result [355]. Let us denote this operator by χ^{mat} . The advantage of using this operator is that it is quite general and can be used for different purposes (see e.g. [88], Chap. ??, Chap. ??). Since the map operator extends a given input tuple by some attribute values, which must be computed by an expression, we need one to express the access to a unique index. For our example, we write

$$\text{IdxAcc}_{\text{dno}}^{\text{Dept}}[y; y.\text{dno} = \text{dno}]$$

to express the lookup of a single (unique) entry in the index on $\text{Dept}.\text{dno}$. We assume that the result is a (pointer to the) tuple containing the key attributes and all data attributes including the TID of some tuple. Then, we have to perform a further attribute access (dereferenciation) if we are interested in only one of the attributes.

Now, we can rewrite the above expression to

$$E_i < E_a > < \chi_{\text{dTID}:(\text{IdxAcc}_{\text{dno}}^{\text{Dept}}[y; y.\text{dno}=\text{dno}])}^{\text{mat}}(\square) >$$

If we further assume that the outer ($E_i < E_a >$) is sorted on dno , then it suffices to remember only the TID for the latest dno . We define the map operator $\chi^{\text{mat},1}$ to do exactly this. A more efficient plan could thus be

$$\text{Sort}_{\text{dno}}(E_i < E_a >) < \chi_{\text{dTID}:(\text{IdxAcc}_{\text{dno}}^{\text{Dept}}[y; y.\text{dno}=\text{dno}])}^{\text{mat},1}(\square) >$$

where, strictly speaking, sorting is not necessary: grouping would suffice.

Consider a general expression of the form $e_1 < e_2 >$. The free variables used in e_2 must be a subset of the variables (attributes) produced by e_1 , i.e. $\mathcal{F}(e_2) \subseteq \mathcal{A}(e_1)$. Even if e_1 does not contain duplicates, the projection of e_1 on $\mathcal{F}(e_2)$ may contain duplicates. If so, materialization could pay off. However, in general, for every binding of the variables $\mathcal{F}(e_2)$, the expression e_2 may produce several tuples. This means that using χ^{mat} is not sufficient. Consider the query

```
select *
from   Emp e, Wine w
where  e.yearOfBirth = w.year
```

If we have no indexes, we can answer this query by a simple join where we only have to decide the join method and which of the relations becomes the outer and which the inner. Assume we have only wines from a few years. (Alternatively, some selection could have been applied.) Then it might make sense to consider the following alternative:

$$\text{Wine}[w] < \sigma_{e.\text{yearOfBirth}=w.\text{year}}(\text{Emp}[e]) >$$

However, the relation **Emp** is scanned once for each **Wine** tuple. Hence, it might make sense to materialize the result of the inner for every **year** value of **Wine** if we have only a few **year** values. In other words, if we have many duplicates for the **year** attribute of **Wine**, materialization may pay off since then we have to scan **Emp** only once for each **year** value of **Wine**. To achieve caching of the inner, in case every binding of its free variables possibly results in many tuples, requires a new operator. Let us call this operator *memox* and denote it by \mathfrak{M} [298, 88]. For the free variables of its only argument, it remembers the set of result tuples produced by its argument expression and does not evaluate it again if it is already cached. Using *memox*, the above plan becomes

$$\text{Wine}[w] < \mathfrak{M}(\sigma_{e.\text{yearOfBirth}=w.\text{year}}(\text{Emp}[e])) >$$

It should be clear that for more complex inners, the *memox* operator can be applied at all branches, giving rise to numerous caching strategies. Analogously to the materializing map operator, we are able to restrict the materialization to the results for a single binding for the free variables if the outer is sorted (or grouped) on the free variables:

$$\text{Sort}_{w.\text{yearOfBirth}}(\text{Wine}[w]) < \mathfrak{M}^1(\sigma_{e.\text{yearOfBirth}=w.\text{year}}(\text{Emp}[e])) >$$

Things can become even more efficient if there is an index on **Emp.yearOfBirth**:

$$\text{Sort}_{w.\text{yearOfBirth}}(\text{Wine}[w]) \\ < \mathfrak{M}^1(\text{Emp}_{\text{yearOfBirth}}[x; x.\text{yearOfBirth} = w.\text{year}] < \chi_{e:*(x.\text{TID}),\mathcal{A}(\text{Emp}):*e}(\square) >) >$$

So far we have seen different operators which materialize values: *Tmp*, \mathfrak{M} , and χ_{mat} . The latter in two variants. As an exercise, the reader is advised to discuss the differences between them.

Assume, we have indexes on both `Emp.yearOfBirth` and `Wine.year`. Besides the possibilities to use either `Emp` or `Wine` as the outer, we now also have the possibility to perform a join on the indexes before accessing the actual `Emp` and `Wine` tuples. Since the index scan produces its output ordered on the key attributes, a simple merge join suffices (and we are back at the latter):

$$\text{Emp}_{\text{yearOfBirth}}[x] \bowtie_{x.\text{yearOfBirth}=y.\text{year}}^{\text{merge}} \text{Wine}_{\text{year}}[y]$$

This example makes clear that the order provided by an index scan can be used to speed up join processing. After evaluating this plan fragment, we have to access the actual `Emp` and `Wine` tuples. We can consider zero, one, or two sorts on their respective tuple identifiers. If the join is sufficiently selective, one of these alternatives may prove more sufficient than the ones we have considered so far.

EX

4.15 Remarks on Access Path Generation

A last kind of optimization we briefly want to mention is *sideways information passing*. Consider a simple join between two relations: $R \bowtie_{R.a=S.b} S$. If we decide to perform a sort merge join or a hash join, we can implement it by first sorting/partitioning R before looking at S . While doing so, we can remember the minimum and maximum value of $R.a$ and use these as a restriction on S such that fewer tuples of S have to be sorted/partitioned. In case we perform a blockwise nested loop join, after the first scan of S we know the minimum and maximum value of $S.b$ and can use these to restrict R .

If the number of distinct values of $R.a$ is small, we could also decide to remember all these values and evaluate perform a semi-join before the actual join. Algebraically, this could be expressed as

$$R \bowtie_{R.a=S.b} (S \bowtie_{S.b=R.a} \Pi_{R.a}(R))$$

An alternative is to use a bitmap to represent the projection of R on a .

The semi-join technique should be well-known from distributed database systems. In deductive database systems, this kind of optimization often carries the attribute *magic*. We will more deeply discuss this issue in Chapter ??.

The following problem is not discussed in the book. Assume that we have fully partitioned a relation vertically into a set of files which are chronologically ordered. Then, the attribute a_i of the j -th tuple can be found at the j -th position of the i -th file. This organization is called *partitioned transposed file* [49]. (Compare this with variant (projection) indexes [551] and small materialized aggregates [527].) The problem is to find an access strategy to all the attribute required by the query given a collection of restriction on some of the relation's attributes. This problem has been discussed in depth by Batory [49]. Full vertical partitioning is also used as the organizing principle of Monet []. Lately, it also gained some interest in the US [].

4.16 Counting the Number of Accesses

4.16.1 Counting the Number of Direct Accesses

After the index scan, we have a set of (distinct) tuple identifiers for which we have to access the original tuples. The question we would like to answer is:

How many pages do we have to read?

Let R be the relation for which we have to retrieve the tuples. Then we use the following abbreviations

N	$ R $	number of tuples in the relation R
m	$\ R\ $	number of pages on which tuples of R are stored
B	N/m	number of tuples per page (<i>blocking factor</i>)
k		number of (distinct) TIDs for which tuples have to be retrieved

We assume that the tuples are uniformly distributed among the m pages. Then, each page stores $B = N/m$ tuples. B is called *blocking factor*.

Let us consider some borderline cases. If $k > N - N/m$ or $m = 1$, then all pages are accessed. If $k = 1$ then exactly one page is accessed. The answer to the general question will be expressed in terms of *buckets* (pages in the above case) and *items* contained therein (tuples in the above case). Later on, we will also use extents, cylinders, or tracks as buckets and tracks or sectors/blocks as items.

We assume that a bucket contains items. The total number of items will be N and the number of requested items will be k . The above question can then be reformulated to how many buckets contain at least one of the k requested items, i.e. how many qualifying buckets exist. We start out by investigating the case where the items are uniformly distributed among the buckets. Two subcases will be distinguished:

1. k distinct items are requested
2. k non-distinct items are requested.

We then discuss the case where the items are non-uniformly distributed.

In any case, the underlying access model is random access. For example, given a tuple identifier, we can directly access the page storing the tuple. Other access models are possible. The one we will subsequently investigate is sequential access where the buckets have to be scanned sequentially in order to find the requested items. After that, we are prepared to develop a model for disk access costs.

Throughout this section, we will further assume that the probability that we request a set with k items is $\frac{1}{\binom{N}{k}}$ for all of the $\binom{N}{k}$ possibilities to select a k -set.⁸ We often make use of established equalities for binomial coefficients. For convenience, the most frequently used equalities are listed in Appendix D.

⁸A k -set is a set with cardinality k .

Selecting k distinct items

Our first theorem was discovered independently by Waters [792] and Yao [833]. We formulate it in terms of buckets containing items. We say a bucket *qualifies* if it contains at least one of the k items we are looking for.

Theorem 4.16.1 (Waters/Yao) *Consider m buckets with n items each. Then there is a total of $N = nm$ items. If we randomly select k distinct items from all items, then the number of qualifying buckets is*

$$\bar{\mathcal{Y}}_n^{N,m}(k) = m * \mathcal{Y}_n^N(k) \quad (4.2)$$

where $\mathcal{Y}_n^N(k)$ is the probability that a bucket contains at least one item. This probability is equal to

$$\mathcal{Y}_n^N(k) = \begin{cases} [1 - p]^k & k \leq N - n \\ 1 & k > N - n \end{cases}$$

where p is the probability that a bucket contains none of the k items. The following alternative expressions can be used to calculate p :

$$p = \frac{\binom{N-n}{k}}{\binom{N}{k}} \quad (4.3)$$

$$= \prod_{i=0}^{k-1} \frac{N - n - i}{N - i} \quad (4.4)$$

$$= \prod_{i=0}^{n-1} \frac{N - k - i}{N - i} \quad (4.5)$$

The second expression (4.4) is due to Yao, the third (4.5) is due to Waters. Palvia and March proved both formulas to be equal [568] (see also [35]). The fraction $m = N/n$ may not be an integer. For these cases, it is advisable to have a Gamma-function based implementation of binomial coefficients at hand (see [597] for details).

Depending on k and n , either the expression of Yao or the one of Waters is faster to compute. After the proof of the above formulas and the discussion of some special cases, we will give several approximations for p .

Proof The total number of possibilities to pick the k items from all N items is $\binom{N}{k}$. The number of possibilities to pick k items from all items not contained in a fixed single bucket is $\binom{N-n}{k}$. Hence, the probability p that a bucket does not qualify is $p = \binom{N-n}{k} / \binom{N}{k}$. Using this result, we can do the following calculation

$$\begin{aligned} p &= \frac{\binom{N-n}{k}}{\binom{N}{k}} \\ &= \frac{(N-n)! \, k!(N-k)!}{k!((N-n)-k)! \, N!} \\ &= \prod_{i=0}^{k-1} \frac{N-n-i}{N-i} \end{aligned}$$

which proves the second expression. The third follows from

$$\begin{aligned}
 p &= \frac{\binom{N-n}{k}}{\binom{N}{k}} \\
 &= \frac{(N-n)! \, k!(N-k)!}{k!((N-n)-k)! \, N!} \\
 &= \frac{(N-n)! \, (N-k)!}{N! \, ((N-k)-n)!} \\
 &= \prod_{i=0}^{n-1} \frac{N-k-i}{N-i}
 \end{aligned}$$

□

Let us list some special cases:

If	then $\mathcal{Y}_m^N(k) =$
$n = 1$	k/N
$n = N$	1
$k = 0$	0
$k = 1$	$B/N = (N/m)N = 1/m$
$k = N$	1

We examine a slight generalization of the first case in more detail. Let N items be distributed over N buckets such that every bucket contains exactly one item. Further let us be interested in a subset of m buckets ($1 \leq m \leq N$). If we pick k items, then the number of buckets within the subset of size m that qualify is

$$m\mathcal{Y}_1^N(k) = m\frac{k}{N} \tag{4.6}$$

In order to see that the two sides are equal, we perform the following calculation:

$$\begin{aligned}
 \mathcal{Y}_1^N(k) &= \left(1 - \frac{\binom{N-1}{k}}{\binom{N}{k}}\right) \\
 &= \left(1 - \frac{\frac{(N-1)!}{k!((N-1)-k)!}}{\frac{N!}{k!(N-k)!}}\right) \\
 &= \left(1 - \frac{(N-1)!k!(N-k)!}{N!k!((N-1)-k)!}\right) \\
 &= \left(1 - \frac{N-k}{N}\right) \\
 &= \left(\frac{N}{N} - \frac{N-k}{N}\right) \\
 &= \frac{N - N + k}{N} \\
 &= \frac{k}{N}
 \end{aligned}$$

Since the computation of $\mathcal{Y}_n^N(k)$ can be quite expensive, several approximations have been developed. The first one was given by Waters [791, 792]:

$$p \approx (1 - k/N)^n$$

This approximation (also described elsewhere [271, 568]) turns out to be pretty good. However, below we will see even better approximations.

For $\bar{\mathcal{Y}}_n^{N,m}(k)$ Whang, Wiederhold, and Sagalowicz gave the following approximation for faster calculation [802]:

$$m * [(1 - (1 - 1/m)^k) + (1/(m^2b) * k(k-1)/2 * (1 - 1/m)^{k-1}) + (1.5/(m^3p^4) * k(k-1)(2k-1)/6 * (1 - 1/m)^{k-1})]$$

A rough estimate is presented by Bernstein, Goodman, Wong, Reeve, and Rothnie [69]:

$$\bar{\mathcal{Y}}_n^{N,m}(k) \approx \begin{cases} k & \text{if } k < \frac{m}{2} \\ \frac{k+m}{2} & \text{if } \frac{m}{2} \leq k < 2m \\ m & \text{if } 2m \leq k \end{cases}$$

An interesting and useful result was derived by Dühr and Saharia [207]. They give two formulas and show that they are lower and upper bounds to Water and Yao's formula. The upper and lower bounds for p are

$$p_{\text{lower}} = \left(1 - \frac{k}{N - \frac{n-1}{2}}\right)^n$$

$$p_{\text{upper}} = \left(\left(1 - \frac{k}{N}\right) * \left(1 - \frac{k}{N - n + 1}\right)\right)^{n/2}$$

for $n = N/m$. Dühr and Saharia claim that the maximal difference resulting from the use of the lower and the upper bound to compute the number of page accesses is 0.224—far less than a single page access.

Selecting k non-distinct items

So far, we assumed that we retrieve k *distinct* items. We could ask the same question for k *non-distinct* items. This question demands a different urn model. In urn model terminology, the former case is an urn model with a *non-replacement* assumption, while the latter case is one with a *replacement* assumption. (Deeper insight into urn models is given by Drmota, Gardy, and Gittenberger [213].)

Before presenting a theorem discovered by Cheung [148], we repeat a theorem from basic combinatorics. We know that the number of subsets of size k of a set with N elements is $\binom{N}{k}$. The following lemma gives us the number of k -multisets⁹ (see, e.g. [718]). The number of k -multisets taken from a set S with $|S|$ elements is denoted by $\left(\left(\begin{matrix} N \\ k \end{matrix}\right)\right)$.

⁹A k -multiset is a multiset with k elements.

Lemma 4.16.2 *Let S be a set with $|S| = N$ elements. Then, the number of multisets with cardinality k containing only elements from S is*

$$\left(\binom{N}{k} \right) = \binom{N+k-1}{k}$$

For a proof we just note that there is a bijection between the k -multisets and the k -subsets of a $N+k-1$ -set. We can go from a multiset to a set by f with $f(\{x_1 \leq \dots \leq x_k\}) = \{x_1+0 < x_2+1 < \dots < x_k+(k-1)\}$ and from a set to a multiset via g with $g(\{x_1 < \dots < x_k\}) = \{x_1-0 < x_2-1 < \dots < x_k-(k-1)\}$.

Theorem 4.16.3 (Cheung) *Consider m buckets with n items each. Then there is a total of $N = nm$ items. If we randomly select k not necessarily distinct items from all items, then the number of qualifying buckets is*

$$\overline{Cheung}_n^{N,m}(k) = m * Cheung_n^N(k) \quad (4.7)$$

where

$$Cheung_n^N(k) = [1 - \tilde{p}] \quad (4.8)$$

with the following equivalent expressions for \tilde{p} :

$$\tilde{p} = \frac{\binom{N-n+k-1}{k}}{\binom{N+k-1}{k}} \quad (4.9)$$

$$= \prod_{i=0}^{k-1} \frac{N-n+i}{N+i} \quad (4.10)$$

$$= \prod_{i=0}^{n-1} \frac{N-1-i}{N-1+k-i} \quad (4.11)$$

Eq. 4.9 follows from the observation that the probability that some bucket does not contain any of the k possibly duplicate items is $\frac{\binom{N-n+k-1}{k}}{\binom{N+k-1}{k}}$. Eq. 4.10 follows from

$$\begin{aligned} \tilde{p} &= \frac{\binom{N-n+k-1}{k}}{\binom{N+k-1}{k}} \\ &= \frac{(N-n+k-1)! \ k!((N+k-1)-k)!}{k!((N-n+k-1)-k)! \ (N+k-1)!} \\ &= \frac{(N-n-1+k)! \ (N-1)!}{(N-n-1)! \ (N-1+k)!} \\ &= \prod_{i=0}^{k-1} \frac{N-n+i}{N+i} \end{aligned}$$

Eq. 4.11 follows from

$$\begin{aligned}
\tilde{p} &= \frac{\binom{N-n+k-1}{k}}{\binom{N+k-1}{k}} \\
&= \frac{(N-n+k-1)! \, k!((N+k-1)-k)!}{k!((N-n+k-1)-k)! \, (N+k-1)!} \\
&= \frac{(N+k-1-n)! \, (N-1)!}{(N+k-1)! \, (N-1-n)!} \\
&= \prod_{i=0}^{n-1} \frac{N-n+i}{N+k-n+i} \\
&= \prod_{i=0}^{n-1} \frac{N-1-i}{N-1+k-i}
\end{aligned}$$

□

Cardenas discovered a formula that can be used to approximate \tilde{p} [102]:

$$(1 - n/N)^k$$

As Cheung pointed out, we can use the theorem to derive the number of distinct items accessed contained in a k -multiset.

Corollary 4.16.4 *Let S be a k -multiset containing elements from an N -set T . Then the number of distinct items contained in S is*

$$\mathcal{D}(N, k) = \frac{Nk}{N+k-1} \quad (4.12)$$

if the elements in T occur with the same probability in S .

We apply the theorem for the special case where every bucket contains exactly one item ($n = 1$). In this case, $\prod_{i=0}^0 \frac{N-1-i}{N-1+k-i} = \frac{N-1}{N-1+k}$. And the number of qualifying buckets is $N(1 - \frac{N-1}{N-1+k}) = N(\frac{N-1+k-N+1}{N-1+k}) = N\frac{k}{N+k-1}$. □

Another way to achieve this formula is the following. There are $\binom{N}{l}$ possibilities to pick l different elements out of the N elements in T . In order to build a k -multiset with l different elements, we must additionally choose $n-l$ elements from the l elements. Thus, we have $\binom{N}{l} \left(\binom{l}{n-l} \right)$ possibilities to build a k -multiset. The total number of multisets is $\left(\binom{N}{l} \right)$. Thus we may conclude that

$$\mathcal{D}(N, k) = l \sum_{l=1}^{\min(N, k)} \frac{\binom{N}{l} \left(\binom{l}{n-l} \right)}{\left(\binom{N}{l} \right)}$$

which can be simplified to the above.

A useful application of this formula is to calculate the size of a projection [148]. Another use is that calculating the number of distinct values contained in a multiset allows us to shift from the model with replacement to a model without replacement. However, there is a difference between

$$\overline{\mathcal{Y}}_n^{N,m}(\text{Distinct}(N, k)) \approx \overline{\text{Cheung}}_n^{N,m}(k)$$

even when computing $\overline{\mathcal{Y}}$ with Eq. 4.5. Nonetheless, for $n \geq 5$, the error is less than two percent. One of the problems when calculating the result of the left-hand side is that the number of distinct items is not necessarily an integer. To solve this problem, we can implement all our formulas using the Gamma-function. But even then a small difference remains.

The approximation given in Theorem 4.16.3 is not too accurate. A better approximation can be calculated from the probability distribution. Denote by $p(\mathcal{D}(N, k) = j)$ the probability that the number of distinct values if we randomly select k items with replacement from N given items equals j . Then

$$p(\mathcal{D}(N, k) = j) = \binom{N}{j} \sum_{l=0}^{j-1} j(-1)^k \binom{j}{l} ((j-l)/N)^k$$

and thus

$$\mathcal{D}(N, k) = \sum_{j=1}^{\min(N,k)} j \binom{N}{j} \sum_{l=0}^{j-1} j(-1)^k \binom{j}{l} ((j-l)/N)^k$$

This formula is quite intense to calculate. We can derive a very good approximation by the following reasoning. We draw k elements from the set T with $|T| = N$ elements. Every element from T can be drawn at most k times. We produce N buckets, one for each element of T . In each bucket, we insert k copies of the according element from T . Then, a sequence of draws from T with duplicates can be represented by a sequence of draws without duplicate by mapping them to different copies. Thus, the first occurrence is mapped to the first element in the according bucket, the second one to the second copy and so on. Then, we can apply formula by Waters and Yao to calculate the number of buckets (and hence elements of T) hit:

$$\mathcal{D}(N, k) = \overline{\mathcal{Y}}_N^{Nk,k}(k)$$

Since the approximation is quite accurate and we already know how to efficiently calculate this formula, this is our method of choice.

Non-Uniform Distribution of Items

In the previous sections, we assumed that

1. every page contains the same number of records, and
2. every record is accessed with the same probability.

We now turn to relax the first assumption. Christodoulakis models the distribution by m numbers n_i (for $1 \leq i \leq m$) if there are m buckets. Each n_i equals the number of records in some bucket i [151]. Luk proposes Zipfian record distribution [492]. However, Ijbema and Blanken say that Water and Yao's formula is still better, as Luk's formula results in too low values [379]. They all come up with the same general formula presented below. Vander Zander, Taylor, and Bitton [843] discuss the problem of correlated attributes which results in some clusteredness. Zahorjan, Bell, and Sevcik discuss the problem where every item is assigned its own access probability [842]. That is, they relax the second assumption. We will come back to these issues in Section ??.

We still assume that every item is accessed with the same probability. However, we relax the first assumption. The following formula derived by Christodoulakis [151], Luk [492], and Ijbema and Blanken [379] is a simple application of Waters's and Yao's formula to a more general case.

Theorem 4.16.5 (Yao/Waters/Christodoulakis) *Assume a set of m buckets. Each bucket contains $n_j > 0$ items ($1 \leq j \leq m$). The total number of items is $N = \sum_{j=1}^m n_j$. If we look up k distinct items, then the probability that bucket j qualifies is*

$$\mathcal{W}_{n_j}^N(k, j) = \left[1 - \frac{\binom{N-n_j}{k}}{\binom{N}{k}}\right] \quad (= \mathcal{Y}_{n_j}^N(k)) \quad (4.13)$$

and the expected number of qualifying buckets is

$$\overline{\mathcal{W}}_{n_j}^{N,m}(k) := \sum_{j=1}^m \mathcal{W}_{n_j}^N(k, j) \quad (4.14)$$

Note that the product formulation in Eq. 4.5 of Theorem 4.16.1 results in a more efficient computation. We make a note of this in the following corollary.

Corollary 4.16.6 *Assume a set of m buckets. Each bucket contains $n_j > 0$ items ($1 \leq j \leq m$). The total number of items is $N = \sum_{j=1}^m n_j$. If we look up k distinct items, then the expected number of qualifying buckets is*

$$\overline{\mathcal{W}}_{n_j}^{N,m}(k) = \sum_{j=1}^m (1 - p_j) \quad (4.15)$$

with

$$p_j = \begin{cases} \prod_{i=0}^{n_j-1} \frac{N-k-i}{N-i} & k \leq n_j \\ 0 & N - n_j < k \leq N \end{cases} \quad (4.16)$$

If we compute the p_j after we have sorted the n_j in ascending order, we can use the fact that

$$p_{j+1} = p_j * \prod_{i=n_j}^{n_{j+1}-1} \frac{N-k-i}{N-i}.$$

We can also use the theorem to calculate the number of qualifying buckets in case the distribution is given by a histogram.

Corollary 4.16.7 For $1 \leq i \leq L$ let there be l_i buckets containing n_i items. Then the total number of buckets is $m = \sum_{i=1}^L l_i$, and the total number of items in all buckets is $N = \sum_{i=1}^L l_i n_i$. For k randomly selected items, the number of qualifying buckets is

$$\overline{W}_{n_j}^{N,m}(k) = \sum_{i=1}^L l_i \mathcal{Y}_{n_j}^N(k) \quad (4.17)$$

Last in this section, let us calculate the probability distribution for the number of qualifying items within a bucket. The probability that $x \leq n_j$ items in a bucket j qualify can be calculated as follows. The number of possibilities to select x items in bucket n_j is $\binom{n_j}{x}$. The number of possibilities to draw the remaining $k - x$ items from the other buckets is $\binom{N-n_j}{k-x}$. The total number of possibilities to distribute k items over the buckets is $\binom{N}{k}$. This shows the following:

Theorem 4.16.8 Assume a set of m buckets. Each bucket contains $n_j > 0$ items ($1 \leq j \leq m$). The total number of items is $N = \sum_{j=1}^m n_j$. If we look up k distinct items, the probability that x items in bucket j qualify is

$$\mathcal{X}_{n_j}^N(k, x) = \frac{\binom{n_j}{x} \binom{N-n_j}{k-x}}{\binom{N}{k}} \quad (4.18)$$

Further, the expected number of qualifying items in bucket j is

$$\overline{\mathcal{X}}_{n_j}^{N,m}(k) = \sum_{x=0}^{\min(k, n_j)} x \mathcal{X}_{n_j}^N(k, x) \quad (4.19)$$

In standard statistics books the probability distribution $\mathcal{X}_{n_j}^N(k, x)$ is called *hypergeometric distribution*.

Let us consider the case where all n_j are equal to n . Then we can calculate the average number of qualifying items in a bucket. With $y := \min(k, n)$ we

have

$$\begin{aligned}
\overline{\mathcal{X}}_{n_j}^{N,m}(k) &= \sum_{x=0}^{\min(k,n)} x \mathcal{X}_n^N(k, x) \\
&= \sum_{x=1}^{\min(k,n)} x \mathcal{X}_n^N(k, x) \\
&= \frac{1}{\binom{N}{k}} \sum_{x=1}^y x \binom{n}{x} \binom{N-n}{k-x} \\
&= \frac{1}{\binom{N}{k}} \sum_{x=1}^y \binom{x}{1} \binom{n}{x} \binom{N-n}{k-x} \\
&= \frac{1}{\binom{N}{k}} \sum_{x=1}^y \binom{n}{1} \binom{n-1}{x-1} \binom{N-n}{k-x} \\
&= \frac{\binom{n}{1}}{\binom{N}{k}} \sum_{x=0}^{y-1} \binom{n-1}{0+x} \binom{N-n}{(k-1)-x} \\
&= \frac{\binom{n}{1}}{\binom{N}{k}} \binom{n-1+N-n}{0+k-1} \\
&= \frac{\binom{n}{1}}{\binom{N}{k}} \binom{N-1}{k-1} \\
&= n \frac{k}{N} = \frac{k}{m}
\end{aligned}$$

Let us consider the even more special case where every bucket contains a single item. That is, $N = m$ and $n_i = 1$. The probability that a bucket contains a qualifying item reduces to

$$\begin{aligned}
\mathcal{X}_1^N(k, x) &= \frac{\binom{1}{x} \binom{N-1}{k-1}}{\binom{N}{k}} \\
&= \frac{\binom{N-1}{k-1}}{\binom{N}{k}} \\
&= \frac{k}{N} \quad (= \frac{k}{m})
\end{aligned}$$

Since x can then only be zero or one, the average number of qualifying items a bucket contains is also $\frac{k}{N}$.

The formulas presented in this section can be used to estimate the number of block/page accesses in case of random direct accesses. As we will see next, other kinds of accesses occur and need different estimates.

4.16.2 Counting the Number of Sequential Accesses

Vector of Bits

When estimating seek costs, we need to calculate the probability distribution for the distance between two subsequent qualifying cylinders. We model the

situation as a bitvector of length B with b bits set to 1. Then B corresponds to the number of cylinders and a 1 indicates that a cylinder qualifies.

Theorem 4.16.9 *Assume a bitvector of length B . Within it b ones are uniformly distributed. The remaining $B - b$ bits are zero. Then the probability distribution of the number j of zeros*

1. *between two consecutive ones,*
2. *before the first one, and*
3. *after the last one*

is given by

$$\mathcal{B}_b^B(j) = \frac{\binom{B-j-1}{b-1}}{\binom{B}{b}} \quad (4.20)$$

A more general theorem (see Theorem 4.16.13) was first presented by Yao [834]. The above formulation is due to Christodoulakis [152].

To see why the formula holds, consider the total number of bitvectors having a one in position i followed by j zeros followed by a one. This number is $\binom{B-j-2}{b-2}$. We can choose $B - j - 1$ positions for i . The total number of bitvectors is $\binom{B}{b}$ and each bitvector has $b - 1$ sequences of the form that a one is followed by a sequence of zeros is followed by a one. Hence,

$$\begin{aligned} \mathcal{B}_b^B(j) &= \frac{(B-j-1)\binom{B-j-2}{b-2}}{(b-1)\binom{B}{b}} \\ &= \frac{\binom{B-j-1}{b-1}}{\binom{B}{b}} \end{aligned}$$

Part 1. of the theorem follows. To prove part 2., we count the number of bitvectors that start with j zeros before the first one. There are $B - j - 1$ positions left for the remaining $b - 1$ ones. Hence, the number of these bitvectors is $\binom{B-j-1}{b-1}$ and part 2 follows. Part 3 follows by symmetry.

We can derive a less expensive way to evaluate the formula for $\mathcal{B}_b^B(j)$ as follows. For $j = 0$, we have $\mathcal{B}_b^B(0) = \frac{b}{B}$. If $j > 0$, then

$$\begin{aligned}
\mathcal{B}_b^B(j) &= \frac{\binom{B-j-1}{b-1}}{\binom{B}{b}} \\
&= \frac{(B-j-1)!}{(b-1)!((B-j-1)-(b-1))!} \\
&= \frac{B!}{b!(B-b)!} \\
&= \frac{(B-j-1)! \ b!(B-b)!}{(b-1)!((B-j-1)-(b-1))! \ B!} \\
&= b \frac{(B-j-1)! \ (B-b)!}{((B-j-1)-(b-1))! \ B!} \\
&= b \frac{(B-j-1)! \ (B-b)!}{(B-j-b)! \ B!} \\
&= \frac{b \ (B-j)! \ (B-b)!}{B-j \ (B-b-j)! \ B!} \\
&= \frac{b}{B-j} \prod_{i=0}^{j-1} \left(1 - \frac{b}{B-i}\right)
\end{aligned}$$

This formula is useful when $\mathcal{B}_b^B(j)$ occurs in sums over j because we can compute the product incrementally.

Corollary 4.16.10 *Using the terminology of Theorem 4.16.9, the expected value for the number of zeros*

1. *before the first one,*
2. *between two successive ones, and*
3. *after the last one*

is

$$\bar{\mathcal{B}}_b^B = \sum_{j=0}^{B-b} j \mathcal{B}_b^B(j) = \frac{B-b}{b+1} \quad (4.21)$$

Let us calculate:

$$\begin{aligned}
\sum_{j=0}^{B-b} j \binom{B-j-1}{b-1} &= \sum_{j=0}^{B-b} (B - (B-j)) \binom{B-j-1}{b-1} \\
&= B \sum_{j=0}^{B-b} \binom{B-j-1}{b-1} - \sum_{j=0}^{B-b} (B-j) \binom{B-j-1}{b-1} \\
&= B \sum_{j=0}^{B-b} \binom{b-1+j}{b-1} - b \sum_{j=0}^{B-b} \binom{B-j}{b} \\
&= B \sum_{j=0}^{B-b} \binom{b-1+j}{j} - b \sum_{j=0}^{B-b} \binom{b+j}{b} \\
&= B \binom{(b-1) + (B-b) + 1}{(b-1) + 1} - b \binom{b + (B-b) + 1}{b+1} \\
&= B \binom{B}{b} - b \binom{B+1}{b+1} \\
&= (B - b \frac{B+1}{b+1}) \binom{B}{b}
\end{aligned}$$

With

$$\begin{aligned}
B - b \frac{B+1}{b+1} &= \frac{B(b+1) - (Bb+b)}{b+1} \\
&= \frac{B-b}{b+1}
\end{aligned}$$

the claim follows.

Corollary 4.16.11 *Using the terminology of Theorem 4.16.9, the expected total number of bits from the first bit to the last one, both included, is*

$$\bar{\mathcal{B}}_{tot}(B, b) = \frac{Bb+b}{b+1} \tag{4.22}$$

To see this, we subtract from B the average expected number of zeros between the last one and the last bit:

$$\begin{aligned}
B - \frac{B-b}{b+1} &= \frac{B(b+1)}{b+1} - \frac{B-b}{b+1} \\
&= \frac{Bb + B - B + b}{b+1} \\
&= \frac{Bb+b}{b+1}
\end{aligned}$$

An early approximation of this formula was discovered by Kollias [435].

Corollary 4.16.12 *Using the terminology of Theorem 4.16.9, the number of bits from the first one and the last one, both included, is*

$$\bar{\mathcal{B}}_{1\text{-span}}(B, b) = \frac{Bb - B + 2b}{b + 1} \quad (4.23)$$

We have two possibilities to argue here. The first subtracts from B the number of zeros at the beginning and the end:

$$\begin{aligned} \bar{\mathcal{B}}_{1\text{-span}}(B, b) &= B - 2\frac{B - b}{b + 1} \\ &= \frac{Bb + B - 2B + 2b}{b + 1} \\ &= \frac{Bb - B + 2b}{b + 1} \end{aligned}$$

The other possibility is to add the number of zeros between the first and the last one and the number of ones:

$$\begin{aligned} \bar{\mathcal{B}}_{1\text{-span}}(B, b) &= (b - 1)\bar{\mathcal{B}}_b^B + b \\ &= (b - 1)\frac{B - b}{b + 1} + \frac{b(b + 1)}{b + 1} \\ &= \frac{Bb - b^2 - B + b + b^2 + b}{b + 1} \\ &= \frac{Bb - B + 2b}{b + 1} \end{aligned}$$

EX or Cor? The number of bits from the first bit to the last one including both ... The distance between the first and the last one ...

Let us have a look at some possible applications of these formulas. If we look up one record in an array of B records and we search sequentially, how many array entries do we have to examine on average if the search is successful?

In [502] we find these formulas used for the following scenario. Let a file consist of B consecutive cylinders. We search for k different keys, all of which occur in the file. These k keys are distributed over b different cylinders. Of course, we can stop as soon as we have found the last key. What is the expected total distance the disk head has to travel if it is placed on the first cylinder of the file at the beginning of the search?

Another interpretation of these formulas can be found in [369, 503]. Assume we have an array consisting of B different entries. We sequentially go through all entries of the array until we have found all the records for b different keys. We assume that the B entries in the array and the b keys are sorted. Further, all b keys occur in the array. On the average, how many comparisons do we need to find all keys?

Vector of Buckets

A more general scenario is as follows. Consider a sequence of m buckets containing n_i items each. Yao [834] developed the following theorem.

Theorem 4.16.13 (Yao) Consider a sequence of m buckets. For $1 \leq i \leq m$, let n_i be the number of items in a bucket i . Then there is a total of $N = \sum_{i=1}^m n_i$ items. Let $t_i = \sum_{l=0}^i n_l$ be the number of items in the first i buckets. If the buckets are searched sequentially, then the number of buckets that have to be examined until k distinct items have been found is

$$C_{n_i}^{N,m}(k, j) = \frac{\binom{t_j}{k} - \binom{t_{j-1}}{k}}{\binom{N}{k}} \quad (4.24)$$

Thus, the expected number of buckets that need to be examined in order to retrieve k distinct items is

$$\bar{C}_{n_i}^{N,m}(k) = \sum_{j=1}^m j C_{n_i}^{N,m}(k, j) = m - \frac{\sum_{j=1}^m \binom{t_{j-1}}{k}}{\binom{N}{k}} \quad (4.25)$$

Applications of this formula can be found in [151, 152, 502, 504, 763]. Manolopoulos and Kollias describe the analogue for the replacement model [502].

Lang, Driscoll, and Jou discovered a general theorem which allows us to estimate the expected number of block accesses for sequential search.

Theorem 4.16.14 (Lang/Driscoll/Jou) Consider a sequence of N items. For a batched search of k items, the expected number of accessed items is

$$A(N, k) = N - \sum_{i=1}^{N-1} \text{Prob}[Y \leq i] \quad (4.26)$$

where Y is a random variable for the last item in the sequence that occurs among the k items searched.

proof?

?

With the help of this theorem, it is quite easy to derive many average sequential accesses for different models.

Cor or EX?

4.16.3 Pointers into the Literature

Segments containing records can be organized differently. Records can be placed randomly in the segment, they can be ordered according to some key, or the segment is organized as a tree. Accordingly, the segment is called random, sequential, or tree-structure. From a segment, records are to be retrieved for a given bag of k keys. The general question then is: how many pages do we have to access? The answer depends on whether we assume the replacement or non-replacement model. Six cases occur. For sequential and tree-structured segments, it also makes sense to distinguish between successful, partially (un-)successful, and (totally) unsuccessful searches. These notions capture the different possibilities where for all, some, none of the k keys records are found. The following table provides some entry points into the literature. It is roughly organized around the above categories. (Remember that we discussed the random file organization at length in Section 4.16.1.)

	non-replacement	replacement
random	[148, 151, 492, 585, 802, 833]	[102, 151, 568, 585]
sequential	[56, 151, 449, 504, 568, 567, 700, 834]	[151, 449, 504, 700]
tree-structured	[449, 448, 504, 567, 589]	[449, 448, 504, 700]

4.17 Disk Drive Costs for N Uniform Accesses

The goal of this section is to derive estimates for the costs (time) for retrieving N cache-missed sectors of a segment S from disk. We assume that the N sectors are read in their physical order on disk. This can be enforced by the DBMS, by the operating system's disk scheduling policy (SCAN policy), or by the disk drive controller.

Remembering the description of disk drives, the total costs can be described as

$$\mathcal{C}_{\text{disk}} = \mathcal{C}_{\text{cmd}} + \mathcal{C}_{\text{seek}} + \mathcal{C}_{\text{settle}} + \mathcal{C}_{\text{rot}} + \mathcal{C}_{\text{headswitch}} \quad (4.27)$$

For brevity, we omitted the parameter N and the parameters describing the segment and the disk drive on which the segment resides. Subsequently, we devote a (sometimes tiny) section to each summand. Before that, we have to calculate the number of qualifying cylinders, tracks, and sectors. These numbers will be used later on.

4.17.1 Number of Qualifying Cylinders, Tracks, and Sectors

If N sectors are to be retrieved, we have to find the number of cylinders qualifying in an extent i . Let S_{sec} denote the total number of sectors our segment contains and $S_{\text{cpe}}(i) = L_i - F_i + 1$ be the number of cylinders of the extent. If the N sectors we want to retrieve are uniformly distributed among the S_{sec} sectors of the segment, the number of cylinders that qualifies in (F_i, L_i, z_i) is $S_{\text{cpe}}(i)$ times 1 minus the probability that a cylinder does not qualify. The probability that a cylinder does not qualify can be computed by deviding the total number of possibilities to chose the N sectors from sectors outside the cylinder by the total number of possibilities to chose N sectors from all S_{sec} sectors of the segment. Hence, the number of qualifying cylinders in the considered extent is:

$$Q_c(i) = S_{\text{cpe}}(i) \mathcal{Y}_{D_{\text{Zspc}}(i)}^{S_{\text{sec}}}(N) = S_{\text{cpe}}(i) \left(1 - \frac{\binom{S_{\text{sec}} - D_{\text{Zspc}}(i)}{N}}{\binom{S_{\text{sec}}}{N}}\right) \quad (4.28)$$

We could also have used Theorem 4.16.13.

Let us also calculate the number of qualifying tracks in a partion i . It can be calculated by $S_{\text{cpe}}(i) D_{\text{tpc}} (1 - \text{Prob}(\mathbf{a\ track\ does\ not\ qualify}))$. The probability that a track does not qualify can be computed by dividing the number of ways to pick N sectors from sectors not belonging to a track divided by the number of possible ways to pick N sectors from all sectors.

$$Q_t(i) = S_{\text{cpe}}(i) D_{\text{tpc}} \mathcal{Y}_{D_{\text{Zspt}}(i)}^{S_{\text{sec}}}(N) = S_{\text{cpe}}(i) D_{\text{tpc}} \left(1 - \frac{\binom{S_{\text{sec}} - D_{\text{Zspt}}(i)}{N}}{\binom{S_{\text{sec}}}{N}}\right) \quad (4.29)$$

Just for fun, we calculate the number of qualifying sectors of an extent in zone i . It can be approximated by

$$Q_s(i) = S_{\text{cpe}}(i) D_{\text{Zspc}}(i) \frac{N}{S_{\text{sec}}} \quad (4.30)$$

Since all $S_{\text{cpe}}(i)$ cylinders are in the same zone, they have the same number of sectors per track, and we could also use Waters/Yao to approximate the number of qualifying cylinders by

$$Q_c(i) = \overline{Y}_{D_{\text{Zspc}}(S_{\text{zone}}(i))}^{S_{\text{cpe}}(i) D_{\text{Zspc}}(S_{\text{zone}}(i)), S_{\text{cpe}}(i)}(Q_s(i)) \quad (4.31)$$

This is a good approximation, as long as $Q_s(i)$ is not too small (e.g. > 4).

4.17.2 Command Costs

The command costs C_{cmd} are easy to compute. Every read of a sector requires the execution of a command. Hence

$$C_{\text{cmd}} = N D_{\text{cmd}}$$

estimates the total command costs.

4.17.3 Seek Costs

We give different alternative possibilities to estimate seek costs. We start with an upper bound by exploring Theorem 4.1.1. The first cylinder we have to visit requires a random seek with cost D_{avgseek} . (Well this does not really give us an upper bound. For a true upper bound we should use $D_{\text{seek}}(D_{\text{cyl}} - 1)$.) After that, we have to visit the remaining $Q_c(i) - 1$ qualifying cylinders. The segment spans a total of $S_{\text{last}}(S_{\text{ext}}) - S_{\text{first}}(1) + 1$ cylinders. Let us assume that the first qualifying cylinder is the first cylinder and the last qualifying cylinder is the last cylinder of the segment. Then applying Theorem 4.1.1 gives us the upper bound


$$C_{\text{seek}}(i) \leq (Q_c(i) - 1) D_{\text{seek}} \left(\frac{S_{\text{last}}(S_{\text{ext}}) - S_{\text{first}}(1) + 1}{Q_c(i) - 1} \right)$$

after we have found the first qualifying cylinder.

We can be a little more precise by splitting the seek costs into two components. The first component C_{seekgap} expresses the costs of finding the first qualifying cylinder and jumping from the last qualifying cylinder of extent i to the first qualifying cylinder of extent $i + 1$. The second component $C_{\text{seekext}}(i)$ calculates the seek costs within an extent i . Figure 4.10 illustrates the situation. The total seek costs then are

$$C_{\text{seek}}(i) = C_{\text{seekgap}} + \sum_{i=1}^{S_{\text{ext}}} C_{\text{seekext}}(i)$$

Since there is no estimate in the literature for C_{seekgap} , we have to calculate it ourselves. After we have done so, we present several alternatives to calculate $C_{\text{seekext}}(i)$.



The upper path illustrates $\mathcal{C}_{\text{seekgap}}$, the lower braces indicate those parts for which $\mathcal{C}_{\text{seekext}}$ is responsible.

Figure 4.10: Illustration of seek cost estimate

The average seek cost for reaching the first qualifying cylinder is D_{avgseek} . How far are we now within the first extent? We use Corollary 4.16.10 to derive that the number of non-qualifying cylinders preceding the first qualifying one in some extent i is

$$\bar{B}_{Q_c(i)}^{S_{cpe}(i)} = \frac{S_{cpe}(i) - Q_c(i)}{Q_c(i) + 1}.$$

The same is found for the number of non-qualifying cylinders following the last qualifying cylinder. Hence, for every gap between the last and the first qualifying cylinder of two extents i and $i + 1$, the disk arm has to travel the distance

$$\Delta_{\text{gap}}(i) := \bar{B}_{Q_c(i)}^{S_{cpe}(i)} + S_{\text{first}}(i + 1) - S_{\text{last}}(i) - 1 + \bar{B}_{Q_c(i+1)}^{S_{cpe}(i+1)}$$

Using this, we get

$$\mathcal{C}_{\text{seekgap}} = D_{\text{avgseek}} + \sum_{i=1}^{S_{\text{ext}}-1} D_{\text{seek}}(\Delta_{\text{gap}}(i))$$

Let us turn to $\mathcal{C}_{\text{seekext}}(i)$. We first need the number of cylinders between the first and the last qualifying cylinder, both included, in extent i . It can be calculated using Corollary 4.16.12:

$$\Xi_{\text{ext}}(i) = \bar{B}_{1\text{-span}}(S_{cpe}(i), Q_c(i))$$

Hence, $\Xi(i)$ is the minimal span of an extent that contains all qualifying cylinders.

Using $\Xi(i)$ and Theorem 4.1.1, we can derive an upper bound for $\mathcal{C}_{\text{seekext}}(i)$:

$$\mathcal{C}_{\text{seekext}}(i) \leq (Q_c(i) - 1) D_{\text{seek}}\left(\frac{\Xi(i)}{Q_c(i) - 1}\right) \quad (4.32)$$

Alternatively, we could formulate this as

$$\mathcal{C}_{\text{seekext}}(i) \leq (Q_c(i) - 1)D_{\text{seek}}(\overline{\mathcal{B}}_{Q_c(i)}^{S_{\text{cpe}}(i)}) \quad (4.33)$$

by applying Corollary 4.16.10.

A seemingly more precise estimate for the expected seek cost within the qualifying cylinders of an extent is derived by using Theorem 4.16.9:

$$\mathcal{C}_{\text{seekext}}(i) = Q_c(i) \sum_{j=0}^{S_{\text{cpe}}(i) - Q_c(i)} D_{\text{seek}}(j + 1) \mathcal{B}_{Q_c(i)}^{S_{\text{cpe}}(i)}(j) \quad (4.34)$$

There are many more estimates for seek times. Older ones rely on a linear disk model but also consider different disk scan policies. A good entry point is the work by Theorey and Pinkerton [756, 757].

4.17.4 Settle Costs

The average settle cost is easy to calculate. For every qualifying cylinder, one head settlement takes place:

$$\mathcal{C}_{\text{settle}}(i) = Q_c(i)D_{\text{rdsettle}} \quad (4.35)$$

4.17.5 Rotational Delay Costs

Let us turn to the rotational delay. For some given track in zone i , we want to read the $Q_t(i)$ qualifying sectors contained in it. On average, we would expect that the read head is ready to start reading in the middle of some sector of a track. If so, we have to wait for $\frac{1}{2}D_{\text{Zscan}}(S_{\text{zone}}(i))$ before the first whole sector occurs under the read head. However, due to track and cylinder skew, this event does not occur after a head switch or a cylinder switch. Instead of being overly precise here, we ignore this half sector pass by time and assume that we EX are always at the beginning of a sector. This is also justified by the fact that we model the head switch time explicitly.

Assume that the head is ready to read at the beginning of some sector of some track. Then, in front of us is a — cyclic, which does not matter — bitvector of qualifying and non-qualifying sectors. We can use Corollary 4.16.11 to estimate the total number of qualifying and non-qualifying sectors that have to pass under the head until all qualifying sectors have been seen. The total rotational delay for the tracks of zone i is

$$\mathcal{C}_{\text{rot}}(i) = Q_t(i) D_{\text{Zscan}}(S_{\text{zone}}(i)) \overline{\mathcal{B}}_{\text{tot}}(D_{\text{Zspt}}(S_{\text{zone}}(i)), Q_{\text{spt}}(i)) \quad (4.36)$$

where $Q_{\text{spt}}(i) = \overline{\mathcal{W}}_1^{S_{\text{sec}}, D_{\text{Zspt}}(S_{\text{zone}}(i))}(N) = D_{\text{Zspt}}(S_{\text{zone}}(i)) \frac{N}{S_{\text{sec}}}$ is the expected number of qualifying sectors per track in extent i . In case $Q_{\text{spt}}(i) < 1$, we set $Q_{\text{spt}}(i) := 1$.

A more precise model is derived as follows. We sum up for all j the product of (1) the probability that j sectors in a track qualify and (2) the average number of sectors that have to be read if j sectors qualify. This gives us the number of

sectors that have to pass the head in order to read all qualifying sectors. We only need to multiply this number by the time to scan a single sector and the number of qualifying tracks. We can estimate (1) using Theorem 4.16.8. For (2) we again use Corollary 4.16.11.

$$\mathcal{C}_{\text{rot}}(i) = S_{\text{cpe}}(i) D_{\text{tpc}} D_{\text{Zscan}}(S_{\text{zone}}(i)) \cdot \sum_{j=1}^{\min(N, D_{\text{Zspt}}(S_{\text{zone}}(i)))} \mathcal{X}_{D_{\text{Zspt}}(S_{\text{zone}}(i))}^{\text{Ssec}}(N, j) \bar{\mathcal{B}}_{\text{tot}}(D_{\text{Zspt}}(S_{\text{zone}}(i)), \beta) \quad (4.37)$$

Another approach is taken by Triantafyllou, Christodoulakis, and Georgiadis [763]. They split the total rotational delay into two components. The first component ($\mathcal{C}_{\text{rotpass}}$) measures the time needed to skip unqualifying sectors and the second ($\mathcal{C}_{\text{rotread}}$) that for (scanning and transferring) the qualifying sectors to the host.

Let us deal with the first component. Assume that j sectors of a track in extent i qualify. The expected position on a track where the head is ready to read is the middle between two qualifying sectors. Since the expected number of sectors between two qualifying sectors is $D_{\text{Zspt}}(S_{\text{zone}}(i))/j$, the expected number of sectors scanned before the first qualifying sector comes under the head is $\frac{D_{\text{Zspt}}(S_{\text{zone}}(i))}{2j}$. The expected positions of j qualifying sectors on the same track is such that the number of non-qualifying sectors between two successively qualifying sectors is the same. Hence, after having read a qualifying sector, $\frac{D_{\text{Zspt}}(S_{\text{zone}}(i))}{j}$ unqualifying sectors must pass by until the next qualifying sector shows up. The total number of unqualifying sectors to be passed if j sectors qualify in a track of zone i is

$$N_s(j, i) = \frac{D_{\text{Zspt}}(S_{\text{zone}}(i))}{2j} + (j-1) \frac{D_{\text{Zspt}}(S_{\text{zone}}(i)) - j}{j} \quad (4.38)$$

Using again Theorem 4.16.8, the expected rotational delay for the unqualifying sectors then is

$$\mathcal{C}_{\text{rotpass}}(i) = S_{\text{cpe}}(i) D_{\text{tpc}} D_{\text{Zscan}}(S_{\text{zone}}(i)) \cdot \sum_{j=1}^{\min(N, D_{\text{Zspt}}(S_{\text{zone}}(i)))} \mathcal{X}_{D_{\text{Zspt}}(S_{\text{zone}}(i))}^{\text{Ssec}}(N, j) N_s(j, i) \quad (4.39)$$

We have to sum up this number for all extents and then add the time needed to scan the N sectors. Hence

$$\mathcal{C}_{\text{rot}} = \sum_{i=1}^{S_{\text{ext}}} \mathcal{C}_{\text{rotpass}}(i) + \mathcal{C}_{\text{rotread}}(i)$$

where the total transfer cost for the qualifying sectors of an extent can be estimated as

$$\mathcal{C}_{\text{rotread}}(i) = Q_s(i) D_{\text{Zscan}}(S_{\text{zone}}(i))$$

4.17.6 Head Switch Costs

The average head switch cost is equal to the average number of head switches that occur times the average head switch cost. The average number of head switch is equal to the number of tracks that qualify minus the number of cylinders that qualify since a head switch does not occur for the first track of each cylinder. Summarizing

$$C_{\text{headswitch}} = \sum_{i=1}^{S_{\text{ext}}} (Q_t(i) - Q_c(i)) D_{\text{hdswitch}} \quad (4.40)$$

where Q_t is the average number of tracks qualifying in an extent.

4.17.7 Discussion

The disk drive cost model derived depends on many parameters. The first bunch of parameters concerns the disk drive itself. These parameters can (and must be) extracted from disk drives by using (micro-) benchmarking techniques [266, 752, 517, 660]. The second bunch of parameters concerns the layout of a segment on disk. The database system is responsible for providing these parameters. The closer it is to the disk, the easier these parameters are extracted. Building a runtime system atop the operating system's file system is obviously a bad idea from the cost model perspective. If instead the storage manager of the runtime system implements cylinder aligned extents (or at least track aligned extents) using a raw I/O interface, the cost model will be easier to develop and much more precise. Again, providing reliable cost models is one of the most important tasks of the runtime system.

We have neglected many problems in our disk access model: partially filled cylinders, pages larger than a block, disk drive's cache, remapping of bad blocks, non-uniformly distributed accesses, clusteredness, and so on. Whereas the first two items are easy to fix, the rest is not so easy. In general, database systems ignore the disk drive cache. The justifying argument is that the database buffer is much larger than the disk drive's cache and, hence, it is very unlikely that we read a page that is not in the database buffer but in the disk cache. However, this argument falls short for non-random accesses. Nevertheless, we will ignore the issue in this book. The interested reader is referred to Shriver's thesis for disk cache modeling [701].

Remapping of bad sectors to other sectors really prevents the development of a precise cost model for disk accesses. Modelling disk drives becomes already a nightmare since a nice partitioning of the disk into zones is no longer possible since some sectors, tracks and even cylinders are reserved for the remapping. So even if no remapping takes place (which is very unlikely), having homogeneous zones of hundreds of cylinders is a dream that will never come true. The result is that we do not have dozens of homogeneous zones but hundreds (if not thousands) of zones of medium homogeneity. These should be reduced to a model of dozens of homogeneous zones such that the error does not become too large. The remaining issues will be discussed later in the book.

EX

There is even more to say about our cost model. A very practical issue arises if the number of qualifying cylinders is small. Then for some extent i , the expected number of qualifying cylinders could be $Q_c(i) = 0.38$. For some of our formulas this is a big problem. As a consequence, special cases for small N , small Q_c , small Q_t have to be developed and implemented.

Another issue is the performance of the cost model itself. The query compiler might evaluate the cost model's formulas thousands or millions of times. Hence, they should be fast to evaluate.

So far, we can adequately model the costs of N disk accesses. Some questions remain. For example, how do we derive the number N of pages we have to access? Do we really need to fetch all N pages from disk or will we find some of them in the buffer? If yes, how many? Further, CPU costs are also an important issue. Deriving a cost model for CPU costs is even more tedious than modelling disk drive costs. The only choice available is to benchmark all parts of a system and then derive a cost model using the extracted parameters. To give examples of parameters to be extracted: we need the CPU costs for accessing a page present in the buffer, for accessing a page absent in the buffer, for a next call of an algebraic operator, for executing an integer addition, and so on. Again, this cannot be done without tools [42, 210, 353, 394, 584].

The bottom line is that a cost model does not have to be accurate, but must lead to correct decisions. In that sense, it must be accurate at the break even points between plan alternatives. Let us illustrate this point by means of our motivating example. If we know that the index returns a single tuple, it is quite likely that the sequential scan is much more expensive. The same might be true for 2, 3, 4, and 5 tuples. Hence, an accurate model for small N is not really necessary. However, as we come close to the costs of a sequential scan, both the cost model for the sequential scan and the one for the index-based access must be correct since the product of their errors is the factor a bad choice is off the best choice. This is a crucial point, since it is easy to underestimate sequential access costs by a factor of 2-3 and overestimate random access cost by a factor of 2-5.

4.18 Concluding Remarks

Learned:

- Open Cost: I/O costs: non-uniform stuff, CPU costs: nothing done
- Wrong cardinality estimates: Open, leads to dynamic qo

4.19 Bibliography

ToDo:

- CPU Costs for B-tree search within inner and leaf pages [446]
- Index/Relations: only joins between building blocks [629]

- RDB/V1: predicate push down (views), 2 phase optimization (local: traditional, global: sharing of tables), five categories for predicates, nested loops evaluation for nested correlated subqueries, use of transitivity of equality, conjunctive normal form, use of min/max value of join column to reduce join cardinality by adding another selection to the other relation ($\min(a) \leq b \leq \max(a)$) for join predicate $a=b$.
- K accesses to unique index: how many page faults if buffer has size b? [642]
- buffer mgmt: [220]
- buffer mgmt: [699]
- buffer mgmt: [450]
- buffer mgmt: [643, 644]
- buffer mgmt: [92]
- buffer mgmt: [150]
- structured, semi-structured, unstructured data: [284] cited in Dono76
- B-trees and their improvements [182]
- Vertical partitioning: [224, 505, 49]
- Horizontal and vertical partitioning: [111]
- set oriented disk access to large complex objects [796, 795], assembly operator: [412],
- large objects: [77, 103, 464]

Part II

Foundations

Chapter 5

Logic, Null, and Boolean Expressions

5.1 Two-valued logic

The Boolean algebra with its operations *not* (\neg), *and* (\wedge), and *or* (\vee) is well-known. The truth tables for these operations is given in Figure 5.1.

5.2 NULL values and three valued logic

Many database management systems (in particular all SQL based relational systems) support a concept of *undefined* or *unknown* values, commonly called *NULL* values. These NULL values represent the fact that an attribute value is unknown, and consequently the result of operations on this attribute can be unknown, too.

As a consequence, NULL values induce a three valued logic, as illustrated by the truth table on the left-hand side of Figure ??: When comparing two values x and y for equality, the result is *true* if both values are not NULL and $x = y$ (using standard equality), *false* if both values are not NULL and not $x = y$, and *unknown* (\perp) if any of the two values is NULL, as then no comparison is possible. These three possible outcomes (*true*, *false*, and \perp) propagate through logical operators, and induce a three-valued logic, as shown in Figure 5.5.

While the three-valued logic correctly captures the uncertainty caused by NULL values, and a result of \perp can be reported back to the user as boolean NULL, it is often necessary to convert three-valued logic results into two-valued logic. For the special case of equality comparisons it is sometimes useful to in-

\neg	$\begin{array}{ c c } \hline \textit{true} & \textit{false} \\ \hline \textit{false} & \textit{true} \\ \hline \end{array}$	\vee	$\begin{array}{ c c c } \hline \textit{true} & \textit{false} & \\ \hline \textit{true} & \textit{true} & \textit{true} \\ \hline \textit{false} & \textit{true} & \textit{false} \\ \hline \end{array}$	\wedge	$\begin{array}{ c c c } \hline \textit{true} & \textit{false} & \\ \hline \textit{true} & \textit{false} & \\ \hline \textit{false} & \textit{false} & \textit{false} \\ \hline \end{array}$
--------	---	--------	--	----------	--

Figure 5.1: Truth tables for two-valued logic

$$\begin{array}{lcl}
 \neg true & \implies & false \\
 \neg false & \implies & true \\
 \\
 p \vee p & \implies & p \\
 p \vee \neg p & \implies & true \\
 p_1 \vee (p_1 \wedge p_2) & \implies & p_1 \\
 p \vee false & \implies & p \\
 p \vee true & \implies & true \\
 \\
 p \wedge p & \implies & p \\
 p \wedge \neg p & \implies & false \\
 p_1 \wedge (p_1 \vee p_2) & \implies & p_1 \\
 p \wedge true & \implies & p \\
 p \vee false & \implies & false
 \end{array}$$

Figure 5.2: Simplification Rules

interpret NULL as a distinct value (for example when eliminating duplicates) and use the NULL aware equality $\stackrel{NULL}{=}$ (right-hand column of Figure 5.4). In general three-valued expressions are usually converted into two-valued expressions by considering \perp as *true-interpreted* or *false-interpreted* (Figure 5.6). An example for false-interpreted \perp values are *where* clauses, which must evaluate to *true* to produce a tuple, while *check* conditions are true-interpreted.

5.3 Simplifying Boolean Expressions

5.4 Optimizing Boolean Expressions

5.5 Bibliography

NULL-values: [637, 469, 638, 470]

Commutativity

$$\begin{array}{ll}
p_1 \vee p_2 & \iff p_2 \vee p_1 & p_1 \wedge p_2 & \iff p_2 \wedge p_1 \\
\exists e_1 \exists e_2 p & \iff \exists e_2 \exists e_1 p & \forall e_1 \forall e_2 p & \iff \forall e_2 \forall e_1 p
\end{array}$$

Associativity

$$(p_1 \vee p_2) \vee p_3 \iff p_1 \vee (p_2 \vee p_3) \qquad (p_1 \wedge p_2) \wedge p_3 \iff p_1 \wedge (p_2 \wedge p_3)$$

Distributivity

$$\begin{array}{ll}
p_1 \vee (p_2 \wedge p_3) & \iff (p_1 \vee p_2) \wedge (p_1 \vee p_3) & p_1 \wedge (p_2 \vee p_3) & \iff (p_1 \wedge p_2) \vee (p_1 \wedge p_3) \\
\exists e (p_1 \vee p_2) & \iff (\exists e p_1) \vee (\exists e p_2) & \forall e (p_1 \wedge p_2) & \iff (\forall e p_1) \wedge (\forall e p_2)
\end{array}$$

Idempotency

$$\begin{array}{ll}
p \vee p & \iff p & p \wedge p & \iff p \\
p \vee \neg p & \iff true & p \wedge \neg p & \iff false \\
p_1 \vee (p_1 \wedge p_2) & \iff p_1 & p_1 \wedge (p_1 \vee p_2) & \iff p_1 \\
p \vee false & \iff p & p \wedge true & \iff p \\
p \vee true & \iff true & p \wedge false & \iff false
\end{array}$$

De Morgan

$$\neg(p_1 \vee p_2) \iff \neg(p_1) \wedge \neg(p_2) \qquad \neg(p_1 \wedge p_2) \iff \neg(p_1) \vee \neg(p_2)$$

Negation of Quantifiers

$$\neg(\forall e p) \iff \exists e(\neg p) \qquad \neg(\exists e p) \iff \forall e(\neg p)$$

Elimination of Negation

$$\neg(\neg(p)) \iff p \qquad \neg t_1 \theta t_2 \implies t_1 \bar{\theta} t_2$$

Conditioned Distributivity($\mathcal{F}(p_1) \cap \mathcal{A}(e) = \emptyset$)

$$\begin{array}{ll}
p_1 \vee (\forall e p_2) & \iff \forall e (p_1 \vee p_2) & p_1 \wedge (\exists e p_2) & \iff \exists e (p_1 \wedge p_2) \\
p_1 \vee (\exists e p_2) & \iff \begin{cases} \exists e(p_1 \vee p_2) & \text{if } e \neq \{\} \\ p_1 & \text{if } e = \{\} \end{cases} \\
p_1 \wedge (\forall e p_2) & \iff \begin{cases} \forall e(p_1 \wedge p_2) & \text{if } e \neq \{\} \\ p_1 & \text{if } e = \{\} \end{cases}
\end{array}$$

Figure 5.3: Laws for two-valued logic

	$x=y$		$x \stackrel{NULL}{=} y$	
	x is <i>NULL</i>	x not <i>NULL</i>	x is <i>NULL</i>	x not <i>NULL</i>
y is <i>NULL</i>	\perp	\perp	<i>true</i>	<i>false</i>
y not <i>NULL</i>	\perp	$x = y$	<i>false</i>	$x = y$

Figure 5.4: Comparison functions in the presence of NULL values

\neg	<i>true</i>	<i>false</i>	\perp	\vee	<i>true</i>	<i>false</i>	\perp	\wedge	<i>true</i>	<i>false</i>	\perp
	<i>false</i>	<i>true</i>	\perp	<i>true</i>	<i>true</i>	<i>true</i>	<i>true</i>	<i>true</i>	<i>true</i>	<i>false</i>	\perp
				<i>false</i>	<i>true</i>	<i>false</i>	\perp	<i>false</i>	<i>false</i>	<i>false</i>	<i>false</i>
				\perp	<i>true</i>	\perp	\perp	\perp	\perp	<i>false</i>	\perp

Figure 5.5: Truth tables for three-valued logic

<i>x</i>	$[x]_{NULL}$	$[x]_{NULL}$
<i>true</i>	<i>true</i>	<i>true</i>
\perp	<i>true</i>	<i>false</i>
<i>false</i>	<i>false</i>	<i>false</i>

Figure 5.6: Interpretation of \perp values

Chapter 6

Functional Dependencies

In many query results attribute values are not independent of each other but have certain dependencies. Keeping track of these dependencies is very useful for many optimizations, for example in the following query

```
select  c.id, n.name
from    customers c, nations n
where   c.nid=n.id
order  by c.id, n.name
```

the *order by* clause can be simplified to *c.id* without affecting the result: *c.id* is the key of *customers*, and thus determines *c.nid*. *c.nid* is joined with *n.id*, which is the key of *nations* and determines *n.name*, thus transitively *c.id* determines *n.name*.

These *functional dependencies* between attributes have been studied primarily in the context of database design, but many optimization steps like order optimization (Chapter 23) and query unnesting (Chapter 14) profit greatly from known functional dependencies. In the following we first study functional dependencies when all attributes are not NULL, then extend this to attributes with NULL values, and finally discuss how functional dependencies are effected by relational operators.

6.1 Functional Dependencies

As illustrated by the previous example, a functional dependency describes how attribute values depend on other attribute values. More formally, a relation R (with $A_1, A_2 \subseteq \mathcal{A}(R)$) satisfies a functional dependency

$$f : A_1 \rightarrow A_2$$

if and only if the following condition holds:

$$\forall t_1, t_2 ([t_1 \in R \wedge t_2 \in R \wedge t_1 \cdot A_1 = t_2 \cdot A_1] \Rightarrow [t_1 \cdot A_2 = t_2 \cdot A_2]).$$

For base relations functional dependencies can be derived from the schema, in particular key constraints and *check* conditions [?]. For intermediate results

additional function dependencies can be induced by algebraic operators, as we will see below.

Once some functional dependencies are known to hold, further functional dependencies can be derived by using Armstrong's axioms [?] (assuming $A_1, A_2, A_3 \subseteq \mathcal{A}(R)$):

1. $A_2 \subseteq A_1 \Rightarrow A_1 \rightarrow A_2$
2. $A_1 \rightarrow A_2 \Rightarrow (A_1 \cup A_3) \rightarrow (A_2 \cup A_3)$
3. $A_1 \rightarrow A_2 \wedge A_2 \rightarrow A_3 \Rightarrow A_1 \rightarrow A_3$

The Armstrong axioms are sound and complete, i.e., it is possible to derive all valid functional dependencies by applying these three axioms. For practical purposes it is often convenient to include three additional rules which can be derived from the original axioms:

4. $A_1 \rightarrow A_2 \wedge A_1 \rightarrow A_3 \Rightarrow A_1 \rightarrow (A_2 \cup A_3)$
5. $A_1 \rightarrow (A_2 \cup A_3) \Rightarrow A_1 \rightarrow A_2 \wedge A_1 \rightarrow A_3$
6. $A_1 \rightarrow A_2 \wedge (A_2 \cup A_4) \rightarrow A_3 \Rightarrow (A_1 \cup A_4) \rightarrow A_3$

Given a set of functional dependencies \mathcal{F} , we denote with \mathcal{F}^+ the *closure* of \mathcal{F} , i.e., the set of all functional dependencies that can be derived from \mathcal{F} by using the inference rules shown above.

Closely related to the concept of functional dependencies is the concept of *keys*: Given a relation R and an attribute set $A \subseteq \mathcal{A}(R)$, A is a *super key* of R if $A \rightarrow \mathcal{A}(R)$ holds in R . Further A is a *key* of R if the following condition holds:

$$\forall A'(A' \subset A \Rightarrow \neg(A' \rightarrow \mathcal{A}(R))).$$

6.2 Functional Dependencies in the presence of NULL values

In the presence of NULL values, a relation R (with $A_1, A_2 \subseteq \mathcal{A}(R)$) satisfies a functional dependency

$$f : A_1 \rightarrow A_2$$

if and only if the following condition holds:

$$\forall t_1, t_2 ([t_1 \in R \wedge t_2 \in R \wedge t_1 \cdot_{A_1} \stackrel{NULL}{=} t_2 \cdot_{A_1}] \Rightarrow [t_1 \cdot_{A_2} \stackrel{NULL}{=} t_2 \cdot_{A_2}]).$$

XXX explain why, discuss lax dependencies

6.3 Deriving Functional Dependencies over algebraic operators

XXX dependency graphs

6.4 Bibliography

Chapter 7

An Algebra for Sets, Bags, and Sequences

7.1 Sets, Bags, and Sequences

7.1.1 Sets

Sets and their laws (see Figure 7.1) should be well-known. A set contains elements drawn from some domain \mathcal{D} . In our case, the domain will often be tuples and we only consider finite sets.

The set operations we are interested in are union (\cup), intersection (\cap), and set difference (\setminus). If the domain consists of tuples, we assume that both arguments have the same schema. That is, the attributes and their domains are the same in both arguments. Otherwise, the expression is not well-typed. In any case, set union and intersection are commutative and associative. Set differ-

$$\begin{array}{llll} X \cup \emptyset & = & X & \\ X \cup X & = & X & \\ X \cup Y & = & Y \cup X & \text{(commutativity)} \\ (X \cup Y) \cup Z & = & X \cup (Y \cup Z) & \text{(associativity)} \\ X \cap \emptyset & = & \emptyset & \\ X \cap x & = & X & \\ X \cap Y & = & y \cap X & \text{(commutativity)} \\ (X \cap Y) \cap Z & = & X \cap (Y \cap z) & \text{(associativity)} \\ X \setminus \emptyset & = & X & \\ \emptyset \setminus X & = & \emptyset & \\ X \setminus Y & \neq & Y \setminus X & \text{(** *wrong **)} \\ (X \setminus Y) \setminus Z & \neq & X \setminus (Y \setminus Z) & \text{(** *wrong **)} \\ X \cup (Y \cap Z) & = & (X \cap Y) \cup (X \cap Z) & \text{(distributivity)} \\ X \cap (Y \cup Z) & = & (X \cup Y) \cap (X \cup Z) & \text{(distributivity)} \\ (X \cup Y) \setminus Z & = & (X \setminus Z) \cup (Y \setminus Z) & \text{(distributivity)} \end{array}$$

Figure 7.1: Laws for Set Operations

ence is neither of them. Expressions containing the empty set can be simplified. Further, some distributivity laws hold.

As we have seen in Chapter 2, algebraic equivalences that reorder algebraic operators form the fundamental core for query optimization. One could discuss the reorderability of each pair of operators resulting in n^2 investigations if the number of operators in the algebra is n . In order to simplify this, we introduce a general argument covering most of the cases. The observation will be that *set-linearity* of set operators implies their reorderability easily.

A unary function f from sets to sets is called *set-linear* (or *homomorph*), if and only if the following two conditions hold for all sets X and Y :

$$\begin{aligned} f(\emptyset) &= \emptyset \\ f(X \cup Y) &= f(X) \cup f(Y) \end{aligned}$$

An n -ary mapping from sets to a set is called *set-linear in its i -th argument*, if and only if for all sets X_1, \dots, X_n and X'_i the following conditions hold:

$$\begin{aligned} f(X_1, \dots, X_{i-1}, \emptyset, X_{i+1}, \dots, X_n) &= \emptyset \\ f(X_1, \dots, X_{i-1}, X_i \cup X'_i, X_{i+1}, \dots, X_n) &= f(X_1, \dots, X_{i-1}, X_i, X_{i+1}, \dots, X_n) \\ &\cup f(X_1, \dots, X_{i-1}, X'_i, X_{i+1}, \dots, X_n) \end{aligned}$$

It is called *set-linear*, if it is linear in all its arguments. For a binary function or operator where we can distinguish between the left and the right argument, we call it left (right) set-linear if it is set-linear in its first (second) argument. Note that if an equivalence with linear mappings on both sides has to be proven, it suffices to prove it for disjoint singleton sets, i.e. sets with one element only.

Set intersection is set-linear, set difference is left set-linear but not right set-linear and set union is neither left nor right set-linear.

A set of elements from a domain \mathcal{D} can be seen as a function from \mathcal{D} to $\{0, 1\}$. For a given set S , this function is called the *characteristic function* of S . It can be defined as $\chi_S(s) = \begin{cases} 0 & \text{if } s \notin S \\ 1 & \text{if } s \in S \end{cases}$. Of course there is a bijection between characteristic functions and sets.

7.1.2 Duplicate Data: Bags

A *bag* or *multiset* can contain every element more than once. It cannot contain an element less than zero times. A typical bag is $\{\{a, b, b\}\}$. Another example is $\{\{a, b\}\}$. The latter bag does not contain any duplicates. Hence, it could also be considered a set.

Hence, for a given bag B , the characteristic function for bags maps every element of a domain D to the set of non-negative integers \mathbb{N}_0 . The characteristic function gives the number of occurrences of each element in the bag. Again, there is a bijection between bags and their characteristic functions. We will only consider finite bags.

The bag union $X \cup Y$ of two bags is defined such that the number of occurrences of some element in the union is the sum of its occurrences in X and Y : The number of occurrences of some element in the bag intersection $X \cap Y$ is the

$$\begin{array}{lll}
X \cup \bar{\emptyset} & = & X \\
X \cup X & = & X \\
X \cup Y & = & Y \cup X \quad (\text{commutativity}) \\
(X \cup Y) \cup Z & = & X \cup (Y \cup Z) \quad (\text{associativity}) \\
X \cap \bar{\emptyset} & = & \bar{\emptyset} \\
X \cap x & = & X \\
X \cap Y & = & Y \cap X \quad (\text{commutativity}) \\
(X \cap Y) \cap Z & = & X \cap (Y \cap Z) \quad (\text{associativity}) \\
X \setminus \bar{\emptyset} & = & X \\
\bar{\emptyset} \setminus X & = & \bar{\emptyset} \\
X \setminus Y & \neq & Y \setminus X \quad (**\textit{wrong}**) \\
(X \setminus Y) \setminus Z & \neq & X \setminus (Y \setminus Z) \quad (**\textit{wrong}**) \\
X \cup (Y \cap Z) & = & (X \cap Y) \cup (X \cap Z) \quad (\text{distributivity}) \\
X \cap (Y \cup Z) & \neq & (X \cup Y) \cap (X \cup Z) \quad (**\textit{wrong}**) \\
(X \cup Y) \setminus Z & \neq & (X \setminus Z) \cup (Y \setminus Z) \quad (**\textit{wrong}**)
\end{array}$$

Figure 7.2: Laws for Bag Operations

minimum of the number its occurrences in X and Y . In the bag difference $X \setminus Y$ the number of occurrences of some element is the difference ($\dot{-}$) of its occurrences in X and Y where $a \dot{-} b$ is defined as $\max(0, a - b)$. Using characteristic functions, we can define

$$\begin{aligned}
\chi_{X \cup Y}(z) &= \chi_X(z) + \chi_Y(z) \\
\chi_{X \cap Y}(z) &= \min(\chi_X(z), \chi_Y(z)) \\
\chi_{X \setminus Y}(z) &= \chi_X(z) \dot{-} \chi_Y(z)
\end{aligned}$$

The laws for sets do not necessarily hold for bags (see Figure 7.2). We will have that bag union and bag intersection are both commutative and associative. Bag difference is neither of them. The only distributivity law that holds for bags is the one that distributes a union over an intersection.

Having every operation twice, once for bags and once for sets is quite inconvenient. Fortunately, for some operations we only need the one for bags. We can get rid of some set operations as follows. Every set can be seen as a bag whose the characteristic function never exceeds one. Let $\bar{I}(S)$ turn a set S into a bag with identical characteristic function. The partial function $\bar{I}^{-1}(B)$ return a bag into a set if the bag's characteristic function does not exceed one. Otherwise let \bar{I}^{-1} be undefined. Let X and Y be two sets. For the intersection function we then have

$$\bar{I}^{-1}(\bar{I}(X) \cup \bar{I}(Y)) = X \cup Y$$

That is, for any two sets X and Y bag union and set union are the same. This gives rise to the notion of *set-faithfulness*. We call a unary function on sets f *set-faithful* if and only if

$$\bar{I}^{-1}(f(\bar{I}(X))) = f(X)$$

holds for all sets X . Analogously, binary functions g are *set-faithful* if and only if

$$\bar{I}^{-1}(g(\bar{I}(X), \bar{I}(Y))) = f(X, Y)$$

holds for all sets X and Y .

$\bar{\setminus}$ and $\bar{\cap}$ are set-faithful. Hence, we can (and often will) simply use \setminus and \cap to denote set or bag difference and intersection. However, $\bar{\cup}$ is not. Hence, we have to carefully distinguish between $\bar{\cup}$ and \cup .

To go from a bag to a set, we have to eliminate duplicates. Let us denote by Π^D the duplicate elimination operation. For a given bag B we then have $\chi_{\Pi^D(B)}(z) = \min(1, \chi_B(z))$.

Instead of working with sets and bags, we can work with bags only by identifying every set S with the bag $\bar{I}(S)$. We can annotate all bags with a property indicating whether it contains duplicates or not. If at some place a set is required and we cannot infer that the bag in that place is duplicate free, we can use Π^D as an enforcer of the set property. Note that for every set S we have $\Pi^D(S) = S$. Hence, Π^D does not do any harm except for the resources it takes. The reasoning whether a given expression produces duplicates or not will be very important, especially in the context of XPath/XQuery.

A unary function f from bags to bags is called *bag-linear* (or *homomorph*), if and only if the following two conditions hold for all bags X and Y :

$$\begin{aligned} f(\bar{\emptyset}) &= \bar{\emptyset} \\ f(X\bar{\cup}Y) &= f(X)\bar{\cup}f(Y) \end{aligned}$$

An n -ary mapping from bags to a bag is called *bag-linear in its i -th argument*, if and only if for all bags X_1, \dots, X_n and X'_i the following conditions hold:

$$\begin{aligned} f(X_1, \dots, X_{i-1}, \bar{\emptyset}, X_{i+1}, \dots, X_n) &= \bar{\emptyset} \\ f(X_1, \dots, X_{i-1}, X_i \cup X'_i, X_{i+1}, \dots, X_n) &= f(X_1, \dots, X_{i-1}, X_i, X_{i+1}, \dots, X_n) \\ &\quad \bar{\cup} f(X_1, \dots, X_{i-1}, X'_i, X_{i+1}, \dots, X_n) \end{aligned}$$

It is called *bag-linear*, if it is linear in all its arguments. For a binary function or operator where we can distinguish between the left and the right argument, we call it left (right) bag-linear if it is bag-linear in its first (second) argument. Note that if an equivalence with linear mappings on both sides has to be proven, it suffices to prove it for disjoint singleton bags, i.e. bags with one element only.

None of bag union, intersection, and difference are left or right bag-linear.

TODO

- $(X\bar{\cup}Y)\bar{\setminus}Z \neq (X\bar{\setminus}Z)\bar{\cup}(Y\bar{\setminus}Z)$
- $\mathcal{A}((R)S) = \mathcal{A}((S)R)$
 $(R\bar{\bowtie}S) \neq (R\bar{\cap}S)$

7.1.3 Ordered Data: Sequences

A sequence is ordered and may contain duplicates. An example sequence is $\langle a, b, b, c, b \rangle$. The length of the sequence is the number of elements it contains.

For any sequence S , the length of the sequence is denoted by $|S|$. The above sequence has length five. The empty sequence (ϵ) contains zero elements and has length zero.

As we consider only finite sequences, a sequence of length $n \geq 0$ has a characteristic function χ from an interval $[0, n[$ to a domain D . Outside $[0, n[$, χ is undefined (\perp). Let S be a sequence. Then $\alpha(S)$ gives us the first element of the sequence, i.e. $\alpha(S) = \chi_S(0)$. For our example sequence, $\alpha(\langle a, b, b, c, b \rangle) = a$. The rest or tail of a sequence S of length n is denoted by $\tau(S)$ and contains all but the first element in the sequence. That is $\chi_{\tau(S)}(i) = \chi_S(i + 1)$. For our example sequence, $\tau(\langle a, b, b, c, b \rangle) = \langle b, b, c, b \rangle$.

Concatenation of two sequences is denoted by \oplus . The characteristic function of the concatenation of two sequences S and T is $\chi_{S \oplus T}(i) = \begin{cases} \chi_S(i) & \text{if } i < |S| \\ \chi_T(i - |S|) & \text{if } i \geq |S| \end{cases}$. As an example, $\langle a, b, b, c, b \rangle \oplus \langle a, b, c \rangle = \langle a, b, b, c, b, a, b, c \rangle$.

We can easily go from a sequence to a bag by just forgetting the order. To convert a bag into a sequence, we typically have to apply a **Sort** operator. In reality however, bags are often represented as (ordered) streams, i.e. they are sequences. This is due to fact that most physical algebras are implemented using the iterator concept introduced in Section 4.6.

Analogously to set and bag linearity, we can introduce *sequence linearity* of unary and n-ary functions on sequences. In the definition, we only have to exchange the set (bag) union operator by concatenation:

A unary function f from sequences to sequences is called *sequence-linear*, if and only if the following two conditions hold for all sequences X and Y :

$$\begin{aligned} f(\epsilon) &= \epsilon \\ f(X \oplus Y) &= f(X) \oplus f(Y) \end{aligned}$$

An n -ary mapping from sequences to a sequence is called *sequence-linear in its i -th argument*, if and only if for all sequences X_1, \dots, X_n and X'_i the following conditions hold:

$$\begin{aligned} f(X_1, \dots, X_{i-1}, \epsilon, X_{i+1}, \dots, X_n) &= \epsilon \\ f(X_1, \dots, X_{i-1}, X_i \cup X'_i, X_{i+1}, \dots, X_n) &= f(X_1, \dots, X_{i-1}, X_i, X_{i+1}, \dots, X_n) \\ &\quad \oplus f(X_1, \dots, X_{i-1}, X'_i, X_{i+1}, \dots, X_n) \end{aligned}$$

It is called *sequence-linear*, if it is sequence-linear in all its arguments. For a binary function or operator where we can distinguish between the left and the right argument, we call it left (right) sequence-linear if it is sequence-linear in its first (second) argument. Note that if an equivalence with linear mappings on both sides has to be proven, it suffices to proof it for disjoint singleton sequences, i.e. sequences with one element only.

7.2 Algebra

Grouping Mapping and Commutativity Comment on $\Pi_A(R_1 \cap R_2) \neq \Pi_A(R_1) \cap \Pi_A(R_2)$ ToDo
 $\Pi_A(R_2) \Pi_A(R_1 \setminus R_2) \neq \Pi_A(R_1) \setminus \Pi_A(R_2)$ ToDo

This section summarizes a logical algebra that is sufficient to express plans for SQL, OQL and XPath/XQuery queries. The algebra is an extension of a union of algebras developed by many people [61, 63, 161, 162, 416, 417, 466, 711]. The most prominent features of the algebra are:

- All operators are polymorphic and can deal with (almost) any kind of complex arguments.
- The operators take arbitrary complex expressions as subscripts. This includes algebraic expressions.
- The algebra is redundant since some special cases of the operators can be implemented much more efficiently. These cases are often introduced as abbreviations.

The remainder of this subsection is organized by operator. The core of the algebra consists of the following operators that are all defined on sets, bags, and sequences. We have already seen the union, intersubsection, and difference operators. The others are selection (σ), projection (Π), map (χ), join (\bowtie), semi-join ($\triangleright\langle$), anti-join (\triangleright), left-outer join ($\bowtie\langle$), d-join ($\langle \cdot \rangle$), map (χ), unnest (μ), and unary (Γ) and binary grouping (Γ). For every algebraic operator we have a subsection discussing it.

While some of the operators are simple extensions of standard operators, others are not widely spread. This forces us to reconsider existing algebraic equivalences as well as to discover new ones. The main question here is reorderability of operators since reorderability is fundamental to any optimization process. Hence, we will discuss reorderability results. Thereby, we will make use of the notion of linearity.

ToDo

Since the join remains the most expensive operator in our algebra, subsection 7.2.11 discusses simplifications of expressions containing joins.

7.2.1 The Operators

The set operators as well as the selection operator and the different join operators (except for the d-join) are known from the relational context. As we will see, the only difference in their definition is that they are extended to express nested queries. For this, we allow the subscripts of these operators to contain full algebraic expressions. Further, to adjust them to the object-oriented context, they do not only deal with relations, but with sets, bags, and sequences of items, where items can be simple tuples, tuples where attributes are allowed to have complex values, and even non-tuple items (e.g. objects or document nodes). This means, that the attribute values are in no way restricted to atomic types but can carry also objects, sets, bags, sequences and so on.

The left-outer join additionally needs some tuning in order to exploit the possibility to have sets as attribute values. For this, it carries a superscript giving a default value for some attribute for those tuples in the left argument for which there is no matching tuple in the right argument. The d-join operation is used for performing a join between two bulk valued items, where the

second one is dependent on the first one. This operator can be used for unnesting nested queries and is in many cases equivalent to a join between two sets with a membership predicate [694]. In some cases (as we will see later on), it corresponds to an unnest operation. We introduced the d-join in order to cope with the values of types that do not have extensions (i.e. there exist no explicit set on which a join could be applied). The d-join is also useful for introducing index structures.

The map operator χ ([416]) is well-known from the functional programming language context. A special case of it, where it adds derived information in form of an added attribute with an according value (e.g. by object-base lookup or by method calls) to each tuple of a set has been proposed in [415, 416]. Later, this special case was given the name *materialization operator* [80].

The unnest operator is known from NF² [654, 636]. It will come in two different flavors allowing us to perform unnesting not only on nested relations but also on attributes whose value is a bulk of elements which are not tuples. The last operator—the *grouping operator*—generalizes the nest operator quite a bit. That is why we renamed it. In fact, there exist two grouping operators, one unary grouping operator and one binary grouping operator. The unary grouping operator groups one set of tuples according to a grouping condition. Further, it can apply an arbitrary expression to the newly formed group. This allows an efficient implementation by saving on intermediate results. The binary grouping operator adds a group to each element in the first argument set. This group is formed from the second argument. The grouping operator will exploit the fact that in the object-oriented context attributes can have set-valued attributes. As we will see, this is useful for both, unnesting nested queries and producing nested results. A variant of the binary grouping operator is also called *nest-join* [720]. Implementations of it are discussed in [120]. There, the binary grouping operator is called *MD-Join*.

7.2.2 Preliminaries

As already mentioned, our algebraic operators can deal not only with standard relations but are polymorphic in the general sense. In order to fix the domain of the operators we need some technical abbreviations and notations. Let us introduce these first.

Since our operators are polymorphic, we need variables for types. We use τ possibly with a subscript to denote types. To express that a certain expression is of type e , we write $e :: \tau$. Starting from concrete names for types and type variables, we can build type expressions the standard way by using type constructors to build tuple types ($[\cdot]$), set types $\{\cdot\}$, bag types $\{\{\cdot\}\}$ and sequence types $\langle \cdot \rangle$. Having two type expressions t_1 and t_2 , we denote by $t_1 \leq t_2$, that t_1 is a subtype of t_2 . It is important to note that this subtype relationship is not based on the sub-/superclass hierarchy found in most object-oriented models. Instead, it simply denotes substitutability. That is the type t_1 provides at least all the attributes and member functions that t_2 provides [101].

Most of our algebraic operators are tuned to work on bulks of tuples. The most important information here is the attributes provided. For this we intro-

duce $\mathcal{A}(A)$. The function \mathcal{A} is defined as follows. $\mathcal{A}(e) = \{a_1, \dots, a_n\}$ if $e :: \{[a_1 : \tau_1, \dots, a_n : \tau_n]\}$, $e :: \{\{[a_1 : \tau_1, \dots, a_n : \tau_n]\}\}$, $e :: \langle [a_1 : \tau_1, \dots, a_n : \tau_n] \rangle$, or $e :: [a_1 : \tau_1, \dots, a_n : \tau_n]$. Giving a set of attributes A , we are sometimes interested in the attributes provided by an expression e which are not in A . For this complement we use the notation $\overline{\mathcal{A}}(e)$ defined as $\mathcal{A}(e) \setminus A$. When e is clear from the context, we use $\overline{\mathcal{A}}$ as a shorthand.

Often, we are not only interested in the set of attributes an expression provides, but also in the set of free variables occurring in an expression e . For this, we introduce $\mathcal{F}(e)$ denoting the set of all free variables of e .

Since the subscripts of our algebraic operators can contain arbitrary expressions, they may contain variables or even free variables. Then, there is a need to get bindings for these variables before the subscript expression can be evaluated. These bindings are taken from the argument(s) of the operator. In order to do so, we need a specified binding mechanism. The λ -notation is such a mechanism and can be used e.g., in case of ambiguities. For our purpose, it suffices if we stick to the following convention.

- For an expression e with free variables $\mathcal{F}(e) = \{a_1, \dots, a_n\}$ and a tuple t with $\mathcal{F}(e) \subseteq \mathcal{A}(t)$ we define $e(t) := e[a_1 \leftarrow t.a_1, \dots, a_n \leftarrow t.a_n]$.¹ Similarly, we define $e(t_1, \dots, t_n)$ for more than a single tuple. Note that the attribute names of the t_i have to be distinct to avoid name conflicts.
- For an expression e with only one free variable x , we define $e(t) = e[x \leftarrow t]$.

The mechanism is very much like the standard binding for the relational algebra. Consider for example a select operation $\sigma_{a=3}(R)$, then we assume that a , the free variable of the subscript expression $a = 3$ is bound to the value of the attribute a of the tuples of the relation R . To express this binding explicitly, we would write for a tuple $t \in R$ $(a = 3)(t)$. Since a is an attribute of R and hence of t , by our convention a is replaced by $t.a$, the value of attribute a of tuple t . Since we want to avoid name conflicts right away, we assume that all variable/attribute names used in a query are distinct. This can be achieved in a renaming step.

Application of a function f to arguments e_i is denoted by either regular (e.g., $f(e_1, \dots, e_n)$) or dot (e.g., $e_1.f(e_2, \dots, e_n)$) notation. The dot notation is used for type associated methods occurring in the object-oriented context.

Last, we introduce the heavily overloaded symbol \circ . It denotes function concatenation and (as a special case) tuple concatenation as well as the concatenation of tuple types to yield a tuple type containing the union of the attributes of the two argument tuple types.

Sometimes it is useful to be able to produce a set/bag/sequence containing only a single tuple with no attributes. This is done by the *singleton scan* operator denoted by \square .

Very often, we are given some database item which is a bulk of other items. Binding these to variables or equivalently, embedding the items into a tuple, we use the notation $e[x]$ for an expression e and a variable/attribute name x . For

¹ $e[v_1 \leftarrow e_1, \dots, v_n \leftarrow e_n]$ denotes a substitution of the variables v_i by the expressions e_i within an expression e .

set valued expressions e , $e[x]$ is defined as $e[x] = \{[x : y] | y \in e\}$. For bags the definition is analogous. For sequence valued expressions e , we define $e[a] = \epsilon$ if e is empty and $e[a] = \langle [a : \alpha(e)] \rangle \oplus \tau(e)[a]$ otherwise.

By id we denote the identity function.

7.2.3 Operator Signatures

We are now ready to define the signatures of the operators of our algebra. Their semantics is defined in a subsequent step. Remember that we consider all operators as being polymorphic. Hence, their signatures are polymorphic and contain type variables, denoted by τ often with an index. We only give the signatures for set operators. For bag operators, just replace the set sign by the

bag sign.

$$\begin{aligned}
\cup & : \{\tau\}, \{\tau\} \rightarrow \{\tau\} \\
\cap & : \{\tau\}, \{\tau\} \rightarrow \{\tau\} \\
\setminus & : \{\tau\}, \{\tau\} \rightarrow \{\tau\} \\
\sigma_p & : \{\tau\} \rightarrow \{\tau\} \\
& \text{if } p : \tau \rightarrow \mathcal{B} \\
\bowtie_p & : \{\tau_1\}, \{\tau_2\} \rightarrow \{\tau_1 \circ \tau_2\} \\
& \text{if } \tau_i \leq \square, p : \tau_1, \tau_2 \rightarrow \mathcal{B}, \text{ and} \\
& \mathcal{A}(\tau_1) \cap \mathcal{A}(\tau_2) = \emptyset \\
\bowtieleft_p & : \{\tau_1\}, \{\tau_2\} \rightarrow \{\tau_1\} \\
& \text{if } \tau_i \leq \square, p : \tau_1, \tau_2 \rightarrow \mathcal{B} \\
\triangleright_p & : \{\tau_1\}, \{\tau_2\} \rightarrow \{\tau_1\} \\
& \text{if } \tau_i \leq \square, p : \tau_1, \tau_2 \rightarrow \mathcal{B} \\
\bowtie_p^{g=c} & : \{\tau_1\}, \{\tau_2\} \rightarrow \{\tau_1 \circ \tau_2^+\} \\
& \text{if } \tau_1 < \square, c :: \tau, \tau_2 < [g : \tau], \\
& \text{+ just adds a null value to the domain of } \tau_2, \\
& \text{if it does not already contain one,} \\
& p : \tau_1, \tau_2 \rightarrow \mathcal{B} \\
\langle \cdot \rangle & : \{\tau_1\} || \{\tau_2\} \rightarrow \{\tau_1 \circ \tau_2\} \\
& \text{if } \tau_i \leq \square, \\
& \mathcal{A}(\tau_1) \cap \mathcal{A}(\tau_2) = \emptyset \\
\chi_f & : \{\tau_1\} \rightarrow \{\tau_2\} \\
& \text{if } f : \tau_1 \rightarrow \tau_2 \\
\Gamma_{g;\theta A;f} & : \{\tau\} \rightarrow \{\tau \circ [g : \tau']\} \\
& \text{if } \tau \leq \square, f : \{\tau\} \rightarrow \tau', A \subseteq \mathcal{A}(\tau) \\
\Gamma_{g;A_1\theta A_2;f} & : \{\tau_1\}, \{\tau_2\} \rightarrow \{\tau_1 \circ [g : \tau']\} \\
& \text{if } \tau_1 \leq \square, f : \{\tau_2\} \rightarrow \tau', A_i \subseteq \mathcal{A}(\tau_i) \text{ for } i = 1, 2 \\
\mu_g & : \{\tau\} \rightarrow \{\tau'\} \\
& \text{if } \tau = [a_1 : \tau_1, \dots, a_n : \tau_n, g : \{\tau_0\}], \\
& \tau_0 \leq \square, \\
& \tau' = [a_1 : \tau_1, \dots, a_n : \tau_n] \circ \tau_0, \\
\mu_{g;c} & : \{\tau\} \rightarrow \{\tau'\} \\
& \text{if } \tau = [a_1 : \tau_1, \dots, a_n : \tau_n, g : \{\tau_0\}], \\
& \tau_0 \not\leq \square, \\
& \tau' = [a_1 : \tau_1, \dots, a_n : \tau_n] \circ [c : \tau_0], \\
\text{flatten} & : \{\{\tau\}\} \rightarrow \{\tau\} \\
\text{max}_{g;m;a\theta;f} & : \{\tau\} \rightarrow [m : \tau_a, g : \tau_f] \\
& \text{if } \tau \leq [a : \tau_a], f : \{\tau\} \rightarrow \tau_f
\end{aligned}$$

7.2.4 Selection

Note that in the following definition there is no restriction on the selection predicate. It may contain path expressions, method calls, nested algebraic operators, etc.

$$\begin{aligned}\sigma_p(e) &= \{x|x \in e, p(x)\} \\ \bar{\sigma}_p(e) &= \{\{x|x \in e, p(x)\}\}\end{aligned}$$

7.2.5 Projection

7.2.6 Map

These operators are fundamental to the algebra. As the operators mainly work on sets of tuples, sets of non-tuples (mostly sets of objects) must be transformed into sets of tuples. This is one purpose of the map operator. Other purposes are dereferenciation, method and function application. Our translation process also pushes all nesting into map operators.

The first definition corresponds to the standard map as defined in, e.g., [416]. The second definition concerns the special case of a materialize operator [415, 80]. The third definition handles the frequent case of constructing a set of tuples with a single attribute out of a given set of (non-tuple) values.

$$\begin{aligned}\chi_{e_2}(e_1) &= \{e_2(x)|x \in e_1\} \\ \chi_{a:e_2}(e_1) &= \{y \circ [a : e_2(y)]|y \in e_1\} \\ e[a] &= \{[a : x]|x \in e\}\end{aligned}$$

Note that the oo map operator obviates the need of a relational projection. Sometimes the map operator is equivalent to a simple projection (or renaming). In these cases, we will use π (or ρ) instead of χ .

7.2.7 Join Operators

The algebra features many different join operators. Besides the complex join predicate, the first four of them are rather standard: join, semi-join, anti-join and left-outer join are defined similarly to their relational counterparts. One difference is that the left-outer join accepts a default value to be given, instead of null, to one attribute of its right argument.

$$\begin{aligned}e_1 \bowtie_p e_2 &= \{y \circ x|y \in e_1, x \in e_2, p(y, x)\} \\ e_1 \bowtie_{<p} e_2 &= \{y|y \in e_1, \exists x \in e_2, p(y, x)\} \\ e_1 \bowtie_{>p} e_2 &= \{y|y \in e_1, \neg \exists x \in e_2 p(y, x)\} \\ e_1 \bowtie_p^{g=c} e_2 &= \{y \circ x|y \in e_1, x \in e_2, p(y, x)\} \cup \\ &\quad \{y \circ z|y \in e_1, \neg \exists x \in e_2 p(y, x), \mathcal{A}(z) = \mathcal{A}(e_2), g \in \mathcal{A}(e_2), \\ &\quad z.g = c, \forall a \in \mathcal{A}(e_2) (z.a \neq NULL \implies a = g)\}\end{aligned}$$

Remember that the function \mathcal{A} used in the last definition returns the set of attributes of a relation.

The last join operator is called *d-join* ($\langle \cdot \rangle$). It is a join between two sets, where the evaluation of the second set may depend on the first set. It is used to translate **from** clauses into the algebra. Here, the range definition of a variable may depend on the value of a formerly defined variable. Whenever possible, d-joins are rewritten into standard joins.

$$e_1 \langle e_2 \rangle = \{y \circ x | y \in e_1, x \in e_2(y)\}$$

Grouping Operators Two grouping operators will be used for unnesting purposes. The first one — called *unary grouping* — is defined on a set and its subscript indicates (i) the attribute receiving the grouped elements (ii) the grouping criterion, and (iii) a function that will be applied to each group.

$$\begin{aligned} \Gamma_{g;\theta A;f}(e) &= \{y.A \circ [g : G] | y \in e, \\ &\quad G = f(\{x | x \in e, x.A \theta y.A\})\} \end{aligned}$$

Note that the traditional nest operator [654] is a special case of unary grouping. It is equivalent to $\Gamma_{g;A=id}$. Note also that the grouping criterion may be defined on several attributes. Then, A and θ represent sequences of attributes and comparators.

The second grouping operator — called *binary grouping* — is defined on two sets. The elements of the second set are grouped according to a criterion that depends on the elements of the first argument.

$$e_1 \Gamma_{g;A_1 \theta A_2;f} e_2 = \{y \circ [g : G] | y \in e_1, G = f(\{x | x \in e_2, y.A_1 \theta x.A_2\})\}$$

In the sequel, the following abbreviations will be used: $\Gamma_{g;A;f}$ for $\Gamma_{g;A=f}$, $\Gamma_{g;A}$ for $\Gamma_{g;A=id}$.

New implementation techniques have to be developed for these grouping operators. Obviously, those used for the nest operator can be used for simple grouping when θ stands for equality. For the other cases, implementations based on sorting seem promising. We also consider adapting algorithms for non-equi joins, e.g. those developed for the band-width join [205]. A very promising approach is the use of θ -tables developed for efficient aggregate processing [163].

Unnest Operators The unnest operators come in two different flavor. The first one is responsible for unnesting a set of tuples on an attribute being a set of tuples itself. The second one unnests sets of tuples on an attribute not being a set of tuples but a set of something else, e.g., integers.

$$\begin{aligned} \mu_g(e) &= \{y.[\mathcal{A}(y) \setminus \{g\}] \circ x | y \in e, x \in y.g\} \\ \mu_{g;c}(e) &= \{y.[\mathcal{A}(y) \setminus \{g\}] \circ [c : x] | y \in e, x \in y.g\} \end{aligned}$$

Flatten Operator The flatten operator flattens a set of sets by unioning the elements of the sets contained in the outer set.

$$flatten(e) = \{y | x \in e, y \in x\}$$

Max Operator The *Max* operator has a very specific use that will be explained in the sequel. Note that an analogous *Min* operator can be defined.

$$Max_{g;m;a\theta;f}(e) = [m : max(\{x.a|x \in e\}), g : f(\{x|x \in e, x.a\theta m\})]$$

This definition is a generalization of the *Max* operator as defined in [161]. Since the equivalences don't care whether we use *Max* or *Min*, we write *Agg* to denote either of them.

Remarks Note that, apart from the χ and *flatten* operations, all these operations are defined on sets of tuples. This guarantees some nice properties among which is the associativity of the join operations. Note also that the operators may take complex expressions in their subscript, therefore allowing nested algebraic expressions. This is the most fundamental feature of the algebra when it comes to express nested queries at the algebraic level. Unnesting is then expressed by algebraic equivalences moving algebraic expression out of the superscript.

The Γ , *flatten* and *Max* operations are mainly needed for optimization purposes, as we will see in the sequel, but do not add power to the algebra. Note that a *Min* operation similar to the *Max* operation can easily be defined.

The algebra is defined on sets whereas most OBMS also manipulate lists and bags. We believe that our approach can easily be extended by considering lists as set of tuples with an added positional attribute and bags as sets of tuples with an added key attribute.

7.2.8 Linearity and Reorderability

Linearity of Algebraic Operators

As already mentioned, the core argument for optimizability of algebraic expressions is the reorderability of their operators. One could discuss the reorderability of each two operators resulting in n^2 investigations if the number of operators in the algebra is n . In order to avoid this, we introduce a general argument covering most of the cases. This argument is that linearity of operators implies their reorderability easily. Hence, let us first look at the linearity property.

A unary mapping $f : \{\tau\} \rightarrow \{\tau'\}$ is called *linear* (or *homomorph*), if and only if

$$\begin{aligned} f(\emptyset) &= \emptyset \\ f(A \cup B) &= f(A) \cup f(B) \end{aligned}$$

If τ , τ' and τ_i are collection types, an n -ary mapping

$$f : \tau_1, \dots, \tau_{i-1}, \{\tau\}, \tau_{i+1}, \dots, \tau_n \rightarrow \{\tau'\}$$

is called *linear in its i -th argument*, iff for all e_1, \dots, e_n, e'_i

$$\begin{aligned} f(e_1, \dots, e_{i-1}, \emptyset, e_{i+1}, \dots, e_n) &= \emptyset \\ f(e_1, \dots, e_{i-1}, e_i \cup e'_i, e_{i+1}, \dots, e_n) &= f(e_1, \dots, e_{i-1}, e_i, e_{i+1}, \dots, e_n) \\ &\quad \cup f(e_1, \dots, e_{i-1}, e'_i, e_{i+1}, \dots, e_n) \end{aligned}$$

It is called *linear*, if it is linear in all its arguments. Note that if an equivalence with linear mappings on both sides has to be proven, it suffices to proof it for disjunct singletons, i.e. sets with one element only.

The following summarizes the findings on the linearity of the algebraic operators:

\cup

\cap

\setminus

π

π^d

χ_f is linear.

$$\begin{aligned}\chi_f(\emptyset) &= \emptyset \\ \chi_f(e_1 \cup e_2) &= \{f(x)|x \in e_1 \cup e_2\} \\ &= \{f(x)|x \in e_1\} \cup \{f(x)|x \in e_2\} \\ &= \chi_f(e_1) \cup \chi_f(e_2)\end{aligned}$$

σ is linear (for a proof see [783])

\bowtie_p is linear (for a proof see [783]). Similarly, $\triangleright<$ is linear. \triangleright is linear in its first argument but obviously not in its second.

\boxtimes is linear in its first argument.

$$\begin{aligned}\emptyset \boxtimes_p^{g=c} e &= \emptyset \\ (e_1 \cup e_2) \boxtimes_p^{g=c} e &= \{y \circ x | y \in e_1 \cup e_2, x \in e, p(y, x)\} \cup \\ &\quad \{y \circ z | y \in e_1 \cup e_2, \neg \exists x \in e p(y, x), \mathcal{A}(z) = \mathcal{A}(e), \\ &\quad \forall a a \neq g \succ z.a = NULL, a.g = c\} \\ &= (e_1 \boxtimes_p^{g=c} e) \cup (e_2 \boxtimes_p^{g=c} e)\end{aligned}$$

To see that \boxtimes is not linear in its second argument consider

$$e_1 \boxtimes_p^{g=c} \emptyset = \emptyset \text{ iff } e_1 = \emptyset$$

$< >$ is linear in its first argument:

$$\begin{aligned}\emptyset < e > &= \emptyset \\ (e_1 \cup e_2) < e > &= \{y \circ x | y \in e_1 \cup e_2, x \in e(y)\} \\ &= \{y \circ x | y \in e_1, x \in e(y)\} \cup \{y \circ x | y \in e_2, x \in e(y)\} \\ &= (e_1 < e >) \cup (e_2 < e >)\end{aligned}$$

Note that the notion of linearity cannot be applied to the second (inner) argument of the d-join, since, in general, it cannot be evaluated independently of the first argument. Below, we summarize some basic observations on the d-join.

$\Gamma_{g;A;f}$ is not linear.

Consider the following counterexample:

$$\begin{aligned} \Gamma_{g;a}(\{[a : 1, b : 1], [a : 1, b : 2]\}) &= \{[a : 1, g : \{[a : 1, b : 1], [a : 1, b : 2]\}]\} \\ &\neq \{[a : 1, g : \{[a : 1, b : 1]\}]\} \cup \{[a : 1, g : \{[a : 1, b : 2]\}]\} \\ &= \Gamma_{g;a}(\{[a : 1, b : 1]\}) \cup \Gamma_{g;a}(\{[a : 1, b : 2]\}) \end{aligned}$$

μ_g is linear.

$$\begin{aligned} \mu_g(\emptyset) &= \emptyset \\ \mu_g(e_1 \cup e_2) &= \{x.[\bar{g}] \circ y \mid x \in e_1 \cup e_2, y \in x.g\} \\ &= \{x.[\bar{g}] \circ y \mid x \in e_1, y \in x.g\} \cup \{x.[\bar{g}] \circ y \mid x \in e_2, y \in x.g\} \\ &= \mu_g(e_1) \cup \mu_g(e_2) \end{aligned}$$

$\mu_{g;c}$ is also linear. This is shown analogously to the linearity of μ_g .

flatten is linear.

$$\begin{aligned} \text{flatten}(e_1 \cup e_2) &= \{x \mid y \in e_1 \cup e_2, x \in y\} \\ &= \{x \mid y \in e_1, x \in y\} \cup \{x \mid y \in e_2, x \in y\} \\ &= \text{flatten}(e_1) \cup \text{flatten}(e_2) \end{aligned}$$

Note that the notion of linearity does not apply to the *max* operator, since it does not return a set.

The concatenation of linear mappings is again a linear mapping. Assume f and g to be linear mappings. Then

$$\begin{aligned} f(g(\emptyset)) &= \emptyset \\ f(g(x \cup y)) &= f(g(x) \cup g(y)) \\ &= f(g(x)) \cup f(g(y)) \end{aligned}$$

Reorderability Laws

From the linearity considerations of the previous subsection, it is easy to derive reorderability laws.

Let $f : \{\tau_1^f\} \rightarrow \{\tau_2^f\}$ and $g : \{\tau_1^g\} \rightarrow \{\tau_2^g\}$ be two linear mappings. If $f(g(\{x\})) = g(f(\{x\}))$ for all singletons $\{x\}$ then

$$f(g(e)) = g(f(e)) \tag{7.1}$$

For the linear algebraic operators working on sets of tuples, we can replace the semantic condition $f(g(\{x\})) = g(f(\{x\}))$ by a set of syntactic criterions. The main issue here is to formalize that two operations do not interfere in their consumer/producer/modifier relationship on attributes. In the relational algebra we have the same problem. Nevertheless, it is often neglected there. Consider for example the algebraic equivalence

$$\sigma_p(R \bowtie S) = (\sigma_p(R)) \bowtie S$$

Then, this algebraic equivalence is true only if the predicate p accesses only those attributes already present in R . Now, for our operators we can be sure that $f(g(e)) = g(f(e))$ for singleton sets e , if g does not consume/produce/modify an attribute that f is going to access and if f is not going to consume/produce/modify an attribute that is accessed by g . In fact, most of the conditions attached to the algebraic equivalences given later on concern this point.

We now consider binary mappings. Let f_1 be a binary mapping being linear in its first argument, f_2 a binary mapping being linear in its second argument, and g a unary linear mapping. If $f_1(g(\{x\}), e') = g(f_1(\{x\}, e'))$ for all x and e' then

$$f_1(g(e), e') = g(f_1(e, e')) \quad (7.2)$$

for all e . Again, we can recast the condition $f_1(g(\{x\}), e') = g(f_1(\{x\}, e'))$ into a syntactical criterion concerning the consumer/producer/modifier relationship of attributes.

Analogously, if $f_2(e, g(\{x\})) = g(f_2(e, \{x\}))$ for all x and e then

$$f_1(e, g(e')) = g(f_1(e, e')) \quad (7.3)$$

for all e' .

Since the outerjoin is not linear in its second argument, we state at least some reorderability results concerning the reordering of joins with outerjoins since much performance can be gained by choosing a (near-) optimal evaluation order. The results are not original but instead taken from [626] and repeated here for convenience:

$$(e_1 \bowtie_{p_{1,2}} e_2) \bowtie_{p_{2,3}} e_3 = e_1 \bowtie_{p_{1,2}} (e_2 \bowtie_{p_{2,3}} e_3) \quad (7.4)$$

$$(e_1 \bowtie_{p_{1,2}} e_2) \bowtie_{p_{2,3}} e_3 = e_1 \bowtie_{p_{1,2}} (e_2 \bowtie_{p_{2,3}} e_3) \quad (7.5)$$

if $p_{2,3}$ is strong w.r.t. e_2

$$(e_1 \bowtie_{p_{1,2}} e_2) \bowtie_{p_{2,3}} e_3 = e_1 \bowtie_{p_{1,2}} (e_2 \bowtie_{p_{2,3}} e_3) \quad (7.6)$$

where an outer join predicate p is strong w.r.t. some expression e_2 , if it yields false if all attributes of the relation to be preserved are NULL.

7.2.9 Reordering of joins and outer-joins

Whereas the join operator is commutative and associative, this is no longer true for outer joins. Especially, joins and outer joins together do not behave

L		
A	B	C
a ₁	b ₁	c ₁
a ₂	b ₂	c ₂

R		
C	D	E
c ₁	d ₁	e ₁
c ₃	d ₂	e ₂

L ⋈ R				
A	B	C	D	E
a ₁	b ₁	c ₁	d ₁	e ₁
a ₂	b ₂	c ₂	-	-

L ⋈ _R R				
A	B	C	D	E
a ₁	b ₁	c ₁	d ₁	e ₁
-	-	c ₃	d ₂	e ₂

L ⋈ _L R				
A	B	C	D	E
a ₁	b ₁	c ₁	d ₁	e ₁
a ₂	b ₂	c ₂	-	-
-	-	c ₃	d ₂	e ₂

Figure 7.3: Outer join examples

associatively. Further, not always are outerjoins and joins reorderable. In this subsection we discuss the reorderability of outerjoins. For a full account on the topic see [626, 254, 263]

The occurrence of an outer join can have several reasons. First, outer joins are part of the SQL 2 specification. Second, outer joins can be introduced during query rewrite. For example, unnesting nested queries or hierarchical views may result in outer joins. Sometimes, it is also possible to rewrite universal quantifiers to outerjoins [766, 187].

The different outer joins can be defined using the *outer union* operator introduced by Codd [170]. Let R_1 and R_2 be two relations and S_1 and S_2 their corresponding attributes. The outerjoin is then defined by padding the union of the relations with *null* values to the schema $S_1 \cup S_2$:

$$R_1 \overset{+}{\cup} R_2 = (R_1 \times \{null_{S_2 \setminus S_1}\}) \cup (R_2 \times \{null_{S_1 \setminus S_2}\}) \tag{7.7}$$

Given this definition of the *outer union* operator, we can define the outer join operations as follows:

$$R_1 \overset{+}{\bowtie}_p R_2 = R_1 \bowtie_p R_2 \overset{+}{\cup} (R_1 \setminus \pi_{S_1}(R_1 \bowtie_p R_2)) \tag{7.8}$$

$$R_1 \overset{+}{\bowtie}_p R_2 = R_1 \bowtie_p R_2 \overset{+}{\cup} (R_1 \setminus \pi_{S_1}(R_1 \bowtie_p R_2)) \overset{+}{\cup} (R_2 \setminus \pi_{S_2}(R_1 \bowtie_p R_2)) \tag{7.9}$$

$$R_1 \bowtie_{p,R} R_2 = R_2 \overset{+}{\bowtie}_p R_1 \tag{7.10}$$

Each outer join preserves the tuples on the according side. For example, the left-outer join preserves the tuples of the relation on the left-hand side. Figure 7.3 gives example applications of the different outer join operations. A *null* value is indicated by a “-”.

Obviously, the left-outer join and the right-outer join are not commutative. To illustrate associativity problems consider the following three relations

R
A
a

S
B C
b -

T
D
d

The results of different left-outer join applications are

R	$\bowtie_{A=B}$	S
A	B	C
a	-	-

S	$\bowtie_{C=D \vee C=null}$	T
B	C	D
b	-	c

(R	$\bowtie_{A=B}$	S)	$\bowtie_{C=D \vee C=null}$	T
A	B	C		D
a	-	-		c

R	$\bowtie_{A=B}$	(S	$\bowtie_{C=D \vee C=null}$	T)
A	B	C		D
a	-	-		-

Hence, in general $(R \bowtie_{p_{RS}} S) \bowtie_{p_{ST}} T \neq r \bowtie_{p_{RS}} (S \bowtie_{p_{ST}} T)$. The problem is that the predicate p_{ST} does not reject *null* values, where a predicate *rejects null values* (or *rejects nulls* for short) in attribute set A , if it evaluates to *false* or *undefined* on every tuple in which all attributes in A are *null*. Using this definition, we have the following identities

$$(R_1 \bowtie_{p_{12}} R_2) \bowtie_{p_{23}} R_3 = R_1 \bowtie_{p_{12}} (R_2 \bowtie_{p_{23}} R_3) \quad (7.11)$$

$$(R_1 \bowtie_{p_{12}} R_2) \bowtie_{p_{23}} R_3 = R_1 \bowtie_{p_{12}} (R_2 \bowtie_{p_{23}} R_3) \\ \text{if } p_{23} \text{ rejects nulls on } \mathcal{A}(R_2) \quad (7.12)$$

$$(R_1 \bowtie_{p_{12}} R_2) \bowtie_{p_{23}} R_3 = R_1 \bowtie_{p_{12}} (R_2 \bowtie_{p_{23}} R_3) \quad (7.13)$$

$$(R_1 \bowtie_{p_{12}} R_2) \bowtie_{p_{23}} R_3 = R_1 \bowtie_{p_{12}} (R_2 \bowtie_{p_{23}} R_3) \\ \text{if } p_{12} \text{ and } p_{23} \text{ reject nulls on } \mathcal{A}(R_2) \quad (7.14)$$

$$(R_1 \bowtie_{p_{12}} R_2) \bowtie_{p_{23}} R_3 = R_1 \bowtie_{p_{12}} (R_2 \bowtie_{p_{23}} R_3) \\ \text{if } p_{23} \text{ rejects nulls on } \mathcal{A}(R_2) \quad (7.15)$$

Further, we can rewrite an outer join to a regular join whenever null-padded tuples are eliminated by some predicate. Equivalences that allow do so and some further ones are given next.

$$R_1 \bowtie_{p_1 \wedge p_2} R_2 = R_1 \bowtie_{p_1} (\sigma_{p_2}(R_2)) \quad \text{if } \mathcal{F}(p_2) \subseteq \mathcal{A}(R_2) \quad (7.16)$$

$$\sigma_{p_1}(R_1 \bowtie_{p_2} R_2) = \sigma_{p_1}(R_1) \bowtie_{p_2} R_2 \quad \text{if } \mathcal{F}(p_1) \subseteq \mathcal{A}(R_1) \quad (7.17)$$

$$\sigma_{p_1}(R_1 \bowtie_{p_2} R_2) = \sigma_{p_1}(R_1 \bowtie_{p_2} R_2) \quad \text{if } p_1 \text{ rejects nulls on } \mathcal{A}(R_2) \quad (7.18)$$

$$\sigma_{p_1}(R_1 \bowtie_{p_2} R_2) = \sigma_{p_1}(R_1 \bowtie_{p_2} R_2) \quad \text{if } p_1 \text{ rejects nulls on } \mathcal{A}(R_2) \quad (7.19)$$

The expression $R_1 \bowtie_{p_{12}} (R_2 \bowtie_{p_{23}} R_3)$ cannot be reordered given the equivalences so far. It is equal to neither $(R_1 \bowtie_{p_{12}} R_2) \bowtie_{p_{23}} R_3$ nor $(R_1 \bowtie_{p_{12}} R_2) \bowtie_{p_{23}} R_3$. In order to allow reorderability on this expression, the *generalized outer join* was introduced by Rosenthal and Galindo-Legaria [626]. It preserves attributes for a set $A \subseteq \mathcal{A}(R_1)$ and is defined as

$$R_1 \bowtie_{p, A} R_2 = (R_1 \bowtie_p R_2) \cup^+ (\pi_A(R_1) \setminus \pi_A(R_1 \bowtie_p R_2)) \quad (7.20)$$

With this definition, we have the following equivalences:

$$\begin{aligned} R_1 \bowtie_{p_{12}}(R_2 \bowtie_{p_{23}} R_3) &= (R_1 \bowtie_{p_{12}} R_2) \bowtie_{p_{23}}^{\mathcal{A}(R_1)} R_3 \quad \text{if } p_{23} \text{ rejects nulls on } \mathcal{A}(R_2) \\ R_1 \bowtie_{p_{12}}(R_2 \bowtie_{p_{23}} R_3) &= (R_1 \bowtie_{p_{12}} R_2) \bowtie_{p_{23}}^{\mathcal{A}(R_1)} R_3 \quad \text{if } p_{23} \text{ rejects nulls on } \mathcal{A}(R_3) \end{aligned} \quad (7.23)$$

The generalized outer join can be generalized to preserve disjoint sets of attributes in order to derive more equivalences [263].

We only gave the basic equivalences for reordering algebraic expressions containing outer joins. General frameworks for dealing with these expressions in toto are presented in [75, 263].

7.2.10 Basic Equivalences for d-Join and Grouping

The d-join and the grouping operators are not linear. Thus, so far, we do not have any reorderability results of these operators. Since they are further quite new, we give some algebraic equivalences which hold despite the fact that they are not linear. This shows that there exist still some kind of optimization which can be performed in the presence of these operators.

Already at the beginning of this subsection, we mentioned that d-join and unnest are closely related. To be more precise, we state the following:

$$e_1 < e_2 > = \mu_g(\chi_{g:e_2}(e_1)) \quad (7.24)$$

$$\pi_{\mathcal{A}(e_2)} e_1 < e_2 > = \mu_g(\chi_{[g:e_2]}(e_1)) \quad (7.25)$$

Between *flatten* and the d-join there also exists a correspondence:

$$\text{flatten}(\chi_{e_2}(e_1)) = \pi_{\mathcal{A}(e_2)}(e_1 < e_2 >) \quad (7.26)$$

The following summarizes basic equivalences on the d-join:

$$e < e > = e \quad (7.27)$$

$$e_1 < e_2 > = e_1 \times e_2 \quad (7.28)$$

$$\text{if } \mathcal{F}(e_2) \cap \mathcal{A}(e_1) = \emptyset$$

$$e_1 < e_2 > = e_1 \bowtie (e_1 < e_2 >) \quad (7.29)$$

$$e_1 < e_2 > < e_3 > = e_1 < e_3 > < e_2 > \quad (7.30)$$

$$\text{if } (\mathcal{A}(e_2) \setminus \mathcal{F}(e_2)) \cap \mathcal{F}(e_3) = \emptyset$$

$$\text{and } (\mathcal{A}(e_3) \setminus \mathcal{F}(e_3)) \cap \mathcal{F}(e_2) \neq \emptyset$$

$$\pi_{\mathcal{A}(e_1)}(e_1 < e_2 >) = \sigma_{e_2 \neq \emptyset}(e_1) \quad (7.31)$$

$$(7.32)$$

We call a function

$$f \text{ extending } \text{ :- } \forall x, y : f(x) \circ y \downarrow \implies f(x \circ y) = f(x) \circ y$$

We call a function

$$f \text{ restricting } \text{ :- } \forall x, y : f(x) \downarrow, x \circ y \downarrow \implies f(x \circ y) = f(x)$$

$\pi_{A':A}$ denotes projection on the attributes A and renaming to A' . Unnesting of operations burried in the d-join can be performed by applying the following equivalence:

$$e_1 < \sigma_p(e_2) > = \sigma_p(e_1 < e_2 >) \quad (7.33)$$

$$e_1 < \sigma_{A_1 \theta A_2}(e_2) > = e_1 \bowtie_{A_1 \theta A_2} e_2 \quad (7.34)$$

if $\mathcal{F}(e_2) \cap \mathcal{A}(e_1) = \emptyset, A_i \subseteq \mathcal{A}(e_i)$

$$\pi_{A':A}(e) < f(\sigma_{A=A'}(e)) > = \mu_g(\pi_{A:A'}(\Gamma_{g;A;f}(e))) \quad (7.35)$$

if $\mathcal{A} \subseteq \mathcal{A}(\cdot)$

$$e_1 < e_2 \bowtie e_3 > = (e_1 < e_2 >) \bowtie e_3 \quad (7.36)$$

if $\mathcal{F}(e_3) \cap \mathcal{A}(e_1) = \emptyset$

$$e_1 < \chi_f(e_2) > = \chi_f(e_1 < e_2 >) \quad (7.37)$$

if f extending

$$\pi_A(e_1 < \chi_f(e_2) >) = \pi_A(\chi_f(e_1 < e_2 >)) \quad (7.38)$$

if $A \subseteq \mathcal{A}(\chi_f(e_2))$, and f restricting

$$e_1 < \mu_g(e_2) > = \mu_g(e_1 < e_2 >) \quad (7.39)$$

$$e_1 < \mu_{g;c}(e_2) > = \mu_{g;c}(e_1 < e_2 >) \quad (7.40)$$

For f_1 being σ or χ and the non-linear unary Γ we still have

$$f_1(\Gamma_{g;=A;f_2}(e)) = \Gamma_{g;=A;f_2}(f_1(e)) \quad (7.41)$$

if $\mathcal{F}(f_1) \cap (A \cup \mathcal{A}(f_2) \cup \{g\}) = \emptyset, (A \cup \mathcal{F}(f_2)) \subseteq \mathcal{A}(f_1(e))$

This equivalence should, e.g., be used from left to right for selections, in order to reduce the cardinality for the Γ operator. For the binary Γ we have

$$f_1(e_1 \Gamma_{g;A_1 \theta A_2; f_2} e_2) = f_1(e_1) \Gamma_{g;A_1 \theta A_2; f_2} e_2 \quad (7.42)$$

since the binary Γ is linear in its first argument.

Lately, work has been reported on the reordering of grouping and join operations despite the fact that grouping is not linear [128, 823, 824, 825, 826]. Since pushing grouping inside join operations can result in a tremendous speed up, let us sketch at least the most basic equivalence:

$$\Gamma_{g;A;agg(e)}(e_1 \bowtie e_2) \equiv e_1 \bowtie (\Gamma_{g;A;agg(e)}(e_2))$$

This sketched equivalence only holds under certain conditions. For details on the conditions and the correctness proof see [825] and Section ??.

7.2.11 Simplifying Expressions Containing Joins

Since the join operation is very expensive, it makes sense to investigate expressions containing joins very intensely in order to discover optimization potential. In this subsection, we do so.

Sometimes, redundant joins can be eliminated:

$$\pi_{\mathcal{A}(e_2)}(e_1 \bowtie_{A_1=A_2} e_2) = e_2 \quad (7.43)$$

$$\text{if } A_1 = \mathcal{A}(e_1), \pi_{A_2}(e_2) \subseteq \pi_{A_2:A_1}(e_1)$$

$$e_1 \bowtie_{A_1=A_{2,3}}^{g:c} (e_2 \bowtie_{A_2=A_3} e_3) = e_1 \bowtie_{A_1=A_3}^{g:c} e_3 \quad (7.44)$$

$$\text{if } A_1 \subseteq \mathcal{A}(e_1), A_2 \subseteq \mathcal{A}(e_2), A_3 \subseteq \mathcal{A}(e_3),$$

$$A_{2,3} \subseteq \mathcal{A}(e_2) \cup \mathcal{A}(e_3), A'_2 = \mathcal{A}(e_1) \cup \mathcal{A}(e_3),$$

$$\pi_{A_2:A'_2}(e_1 \bowtie_{A_1=A_3} e_3) \subseteq e_2$$

$$e_1 \Gamma_{g;A_1=A_{2,3};f}(e_2 \bowtie_{A_2=A_3} e_3) = e_1 \Gamma_{g;A_1=A_2;f} e_3 \quad (7.45)$$

$$\text{if } A_1 \subseteq \mathcal{A}(e_1), A_2 \subseteq \mathcal{A}(e_2), A_3 \subseteq \mathcal{A}(e_3),$$

$$A_{2,3} \subseteq \mathcal{A}(e_2) \cup \mathcal{A}(e_3), A'_2 = \mathcal{A}(e_1) \cup \mathcal{A}(e_3),$$

$$\pi_{A_2:A'_2}(e_1 \bowtie_{A_1=A_3} e_3) \subseteq e_2$$

Equivalences 7.44 and 7.45 bear similarity to equivalence EJA of [783] (p. 107), where the outer-aggregation is used instead of semi-join and binary grouping, respectively. Note that for checking conditions of the form $\pi_{A_2}(e_2) \subseteq \pi_{A_2:A_1}(e_1)$ subtyping implying subsetrelationships on type extensions plays a major role in object bases.

The following shows how to turn an outerjoin into a join:

$$e_1 \bowtie_p^{g:c} e_2 = e_1 \bowtie_p e_2 \quad (7.46)$$

$$\text{if } \neg p(c)$$

$$\sigma_{p_s}(e_1 \bowtie_{p_j}^{g:c} e_2) = \sigma_{p_s}(e_1 \bowtie_{p_j} e_2) \quad (7.47)$$

$$\text{if } \neg p_s(c)$$

$$\sigma_{f_1 \theta f_2}(e_1 \bowtie_{p_j}^{g:c} e_2) = \sigma_{f_1 \theta f_2}(e_1 \bowtie_{p_j} e_2) \quad (7.48)$$

$$\text{if } \neg(f_1 \theta f_2)(c)$$

We can easily check these conditions if some predicate $(f_1 \theta f_2)(c)$ is—after inserting c for the free variable—, $i\theta 0$ with i constant, or $f_1 \in \emptyset$, or a similar simple form. An application of these equivalences can sometimes be followed by a removal of the join.

Note, that the latter equivalence depends on the knowledge we have on the selection predicate. Note also, that the special outerjoin is only introduced by unnesting nested χ operations. Hence, we could combine the equivalences introducing the outerjoin and replacing it by a join into one. In case we have even more information on the selection predicate than above, more specifically, if it depends on a *max* or *min* aggregate, we can do so in a very efficient way:

$$\sigma_{a=m}(e_1)[m : \text{agg}(\chi_b(e_2))] = \text{Agg}_{g;m;a}(e_1).g \quad (7.49)$$

$$\text{if } \pi_a(e_1) = \pi_b(e_2)$$

$$\chi_{g;\sigma_{a=m}(e_2)}(\chi_{m:\text{agg}(e_1)}(e)) = \chi_{\text{Agg}_{g;m;a}(e_1)}(e) \quad (7.50)$$

$$\text{if } \pi_a(e_1) = \pi_b(e_2)$$

7.2.12 Reordering Joins and Grouping

Introduction

In general, join and grouping operations are not reorderable. Consider the following relations R and S

R	A	B
	a	5
	a	6

S	A	C
	a	7
	a	8

Joining these relations R and S results in

$R \bowtie S$	A	B	C
	a	5	7
	a	5	8
	a	6	7
	a	6	8

Applying $\Gamma_{A;count(*)}$ to R and $R \bowtie S$ yields

$\Gamma_{A;count(*)}(R)$	A	count (*)
	a	2

$\Gamma_{A;count(*)}(R \bowtie S)$	A	count (*)
	a	4

Compare this to the result of $\Gamma_{A;count(*)}(R) \bowtie S$:

$\Gamma_{A;count(*)}(R) \bowtie S$	A	count (*)	C
	a	2	7
	a	2	8

Hence $\Gamma_{A;count(*)}(R) \bowtie S \neq \Gamma_{A;count(*)}(R \bowtie S)$.

Since grouping and join operations are in general not reorderable, it is important that a query language determines the order of grouping and join operators properly. In SQL, the grouping operator is applied after the join operators of a query block.

For example, given the relations schemata

Emp (eid, name, age, salary) and
Sold (sid, eid, date, product_id, price)

and the query

```

select    e.eid, sum (s.price) as amount
from      Emp e, Sold s
where     e.eid = s.eid and
           s.date between "2000-01-01" and "2000-12-31"
group by s.eid, s.name

```

results in the algebraic expression

$$\Pi_{e.eid, amount} (\Gamma_{s.eid; amount: \text{sum}(s.price)} (Emp[e] \bowtie_{e.eid=s.eid} \sigma_p (Sold[s])))$$

where p denotes

$$s.date \geq '2000 - 01 - 01' \wedge s.date \leq '2000 - 12 - 31'$$

Figure 20.1 (a) shows this plan graphically. Note that the grouping operator is executed last in the plan.

Now consider the plan where we push the grouping operator down:

$$\Pi_{e.eid,amount} (Emp[e] \bowtie_{e.eid=s.eid} (\Gamma_{s.eid,amount:\mathbf{sum}(s.price)} (\sigma_p (Sold[s]))))$$

This plan (see also Figure 20.1 (b)) is equivalent to the former plan. Moreover, if the grouping operator strongly reduces the cardinality of

$$\sigma_{s.date \geq \dots} (Sold[s])$$

because every employee sells many items, then the latter plan might become cheaper since the join inputs are smaller than in the former plan. This motivates the search for conditions under which join and grouping operators can be reordered. Several papers discuss this reorderability [129, 823, 824, 825, 826]. We will summarize their results in subsequent subsections.

Lazy and eager group by

Lazy group by pulls a group operator up over a join operator [823, 824, 825, 826]. Eager group by does the opposite.

Consider the query:

select [all distinct]	$A, \overrightarrow{F}(B)$
from	R, S
where	$p_R \wedge p_S \wedge p_{R,S}$
group by	G

with

$$G = G_R \cup G_S, G_R \subseteq \mathcal{A}(R), G_S \subseteq \mathcal{A}(S),$$

$$\mathcal{F}(p_R) \subseteq \mathcal{A}(R), \mathcal{F}(p_S) \subseteq \mathcal{A}(S)$$

$$\mathcal{F}(p_{R,S}) \subseteq \mathcal{A}(R) \cup \mathcal{A}(S)$$

$$B \subseteq \mathcal{A}(R) \quad A = A_R \cup A_S, A_R \subseteq G_R, A_S \subseteq G_S$$

$$\alpha_R = G_R \cup \mathcal{F}(p_{R,S}) \setminus \mathcal{A}(S) \quad \kappa_R \text{ key of } R$$

```

      project[e.eid, amount]

      group by[e.eid; amount:sum(price)]

      join[e.eid=s.eid]

Emp[e]          select[s.date between ...]

                Sold[s]
(a)

```

```

      project[e.eid, amount]

      join[e.eid = s.eid]

Emp[e]          group by[e.eid, amount:sum(price)]

                select[s.date ...]

                Sold[s]
(b)

```

Figure 7.4: Two equivalent plans

$$\alpha_S = G_S \cup \mathcal{F}(p_{R,S}) \setminus \mathcal{A}(R) \quad \kappa_S \text{ key of } S$$

We are interested in the conditions under which the query can be rewritten into

```

select[all | distinct]  $A, FB$ 
from  $R', S'$ 
where  $p_{R,S}$ 

with  $R'(\alpha_R, FB) \equiv$ 

select all  $\alpha_R, \overrightarrow{F}(B)$  as  $FB$ 

```

	from	R
	where	p_R
	group by	α_R

and

	$S'(\alpha_S) \equiv$
--	-----------------------

select all	α_R
from	S
where	p_S

The following equivalence expresses this rewrite in algebraic terms.

$$\begin{aligned} & \Pi_{A,F}^{[d]} \left(\Gamma_{G;F:\vec{F}(B)} \left(\sigma_{p_R}(R) \bowtie_{p_{R,S}} \sigma_{p_S}(S) \right) \right) \equiv \\ & \Pi_{A,F}^{[d]} \left(\Gamma_{\alpha_R;F:\vec{F}(B)} \left(\sigma_{p_R}(R) \right) \bowtie_{p_{R,S}} \sigma_{p_S}(S) \right) \end{aligned}$$

holds iff in $\sigma_{p_R \wedge p_S \wedge p_{R,S}}(R \times S)$ the following functional dependencies hold:

$FD_1 : G \rightarrow \alpha_R$

$FD_2 : \alpha_R, G_S \rightarrow \kappa_S$

Note that since $G_S \subseteq G$, this implies $G \rightarrow \kappa_S$.

FD_2 implies that for any group there is at most one join partner in S . Hence, each tuple in $\Gamma_{\alpha_R;F:\vec{F}(B)}(\sigma_{p_R}(R))$ contributes at most one row to the overall result.

FD_1 ensures that each group of the expression on the left-hand side corresponds to at most one group of the group expression on the right-hand side.

We now consider queries with a **having** clause.

In addition to the assumptions above, we have that the tables in the **from** clause can be partitioned into R and S such that R contains all aggregated columns of both the **select** and the **having** clause. We further assume that conjunctive terms in the **having** clause that do not contain aggregate functions have been moved to the **where** clause.

Let the predicate of the **having** clause have the form $H_R \wedge H_0$ where $H_R \subseteq \mathcal{A}(R)$ and $H_0 \subseteq R \cup S$ where H_0 only contains non-aggregated columns from S .

We now consider all queries of the form

select [all distinct]	$A, \vec{F}(B)$
from	R, S
where	$p_R \wedge p_S \wedge p_{R,S}$
group by	G
having	$H_0 \left(\vec{F}_0(B) \right) \wedge H_R \left(\vec{F}_R(B) \right)$

where \vec{F}_0 and \vec{F}_R are vectors of aggregate functions on the aggregated columns B .

An alternative way to express such a query is

select [all distinct]	G, FB
from	R', S
where	$c_S \wedge c_{R,S} \wedge H_0(F_0B)$
where	$R'(\alpha_R, FB, F_0B) \equiv$
select all	$\alpha_R, \vec{F}(B) \text{ as } FB, \vec{F}_0(B) \text{ as } F_0B$
from	R
where	c_R
group by	α_R
having	$H_R(\vec{F}_R(B))$

The according equivalence is [825]:

$$\begin{aligned} & \Pi_{G,F} \left(\sigma_{H_R \wedge H_0} \left(\Gamma_{G;F:\vec{F}(B),F_R:\vec{F}_R(B),F_0:\vec{F}_0(B)} \left(\sigma_{p_R \wedge p_S \wedge p_{R,S}} (R \times S) \right) \right) \right) \\ & \equiv \\ & \Pi_{G,F} \left(\sigma_{p_R, S \wedge p_S \wedge H_0(F_0)} \right) \left(\Pi_{G,F,F_0} \left(\sigma_{H_R} \left(\Gamma_{G;F:\vec{F}(B),F_R:\vec{F}_R(B),F_0:\vec{F}_0(B)} (R) \right) \right) \times S \right) \end{aligned}$$

Coalescing Grouping

In this subsection we introduce *coalescing grouping* which slightly generalizes *simple coalescing grouping* as introduced in [129].

We first illustrate the main idea by means of an example.

Given two relation schemes

Sales (pid, deptid, total_price)
Department (did, name, region)

the query

select region, **sum** (total_price) as s
from Sales, Department
where did = deptid
group by region

is straightforwardly translated into the following algebraic expression:

$$\Gamma_{region;s:sum(total_price)}(\text{Sales} \bowtie_{deptid=did} \text{Department})$$

Note that Equivalence ?? cannot be applied here. However, if there are many sales performed by a department, it might be worth reducing the cardinality of the left join input by introducing an additional group operator. The result is

$$\Gamma_{region;s=sum(s')} \left(\Gamma_{deptid;s':sum(total_price)}(\text{Sales}) \bowtie_{deptid=did} \text{Department} \right)$$

Note that we must keep the outer grouping.

That is, we introduced an additional group operator to reduce the cardinality of sales. This way, all subsequent joins (only one in this case) become cheaper and the additional group operator may result in a better plan.

We have the following restrictions for this subsection:

1. There are no NULL-values allowed for attributes occurring in the query.
2. All queries are of the form **select all**.
That is **select distinct** is not allowed.
3. All aggregate functions agg must fulfill $\text{agg}s_1 \cup s_2 = \text{agg}\{\text{agg}(s_1), \text{agg}(s_2)\}$ for bags s_1 and s_2 .
This has two consequences:
 - Allowed are only sum, min, max. Not allowed are avg and count.
 - For any allowed aggregate function we only allow for **agg(all ...)**.
Forbidden is **agg(distinct ...)**.
4. The query is a single-block conjunctive query with no **having** and no **order by** clause.

The above transformation is an application of the following equivalence, where R_1 and R_2 can be arbitrary algebraic expressions:

$$\Gamma_{G;A}(R_1 \bowtie_p R_2) \equiv \Gamma_{G;A_2}(\Gamma_{G;A_1}(R_1) \bowtie_p R_2) \quad (7.51)$$

with

$$\begin{aligned} A &= A_1 : \text{agg}_1(e_1), \dots, A_n : \text{agg}_n(e_n) \\ A_1 &= A_1^1 : \text{agg}_1^1(e_1), \dots, A_n^1 : \text{agg}_n^1(e_n) \\ A_2 &= A_1 : \text{agg}_1^2(A_1^1), \dots, A_n : \text{agg}_n^2(A_n^1) \\ G_1 &= (\mathcal{F}(p) \cup G) \cap \mathcal{A}(R_1) \end{aligned}$$

Further, the following condition must hold for all $i(1 \leq i \leq n)$:

$$\text{agg}_i \left(\bigcup_k S_k \right) = \text{agg}_i^2 \left(\bigcup_k \{\text{agg}_i^1(S_i)\} \right)$$

In the above example, we had $\text{agg}_1 = \text{agg}_1^1 = \text{agg}_1^2 = \mathbf{sum}$.

We now prove the correctness of Equivalence 20.1.

Proof:

First, note that

$$R_1 \bowtie_p R_2 = \bigcup_{t_2 \in R_2} R_1 \bowtie_p \{t_2\} \quad (7.52)$$

Second, note that for a given t_2

$$\begin{aligned} \Gamma_{G;A}(R_1[t_1]) \bowtie_p \{t_2\} &= \sigma_{p(t_1 \circ t_2)}(\Gamma_{G;A}(R_1[t_1])) \\ &= \Gamma_{G;A}(\sigma_{p(t_1 \circ t_2)}(R_1[t_1])) \\ &= \Gamma_{G;A}(R_1[t_1] \bowtie_p \{t_2\}) \end{aligned} \quad (7.53)$$

holds where we have been a little sloppy with t_1 . Applying (20.2) and (20.3) to $\Gamma_{G_1;A_1}(R_1) \bowtie_p R_2$, the inner part of the right-hand side of the equivalence yields:

$$\begin{aligned} \Gamma_{G_1;A_1}(R_1) \bowtie_p R_2 &= \bigcup_{t_2 \in R_2} \Gamma_{G_1;A_1}(R_1) \bowtie_p \{t_2\} \\ &= \bigcup_{t_2 \in R_2} \Gamma_{G_1;A_1}(R_1 \bowtie_p \{t_2\}) \end{aligned} \quad (7.54)$$

Call the last expression X.

Then the right-hand side of our equivalence becomes

$$\begin{aligned} \Gamma_{G;A_2}(X) &= \{t \circ a_2 \mid t \in \Pi_G(X), a_2 = (A_1 : a_1^2, \dots, A_n : a_n^2), \\ &\quad a_i^2 = \text{agg}_i^2(\{s.A_i^1 \mid s \in X, S|_G = t\})\} \end{aligned} \quad (7.55)$$

Applying (20.2) to the left-hand side of the equivalence yields:

$$\Gamma_{G;A}(R_1 \bowtie_p R_2) = \Gamma_{G;A} \left(\bigcup_{t_2 \in R_2} \{t_1 \circ t_2 \mid t_1 \in R_1, p(t_1 \circ t_2)\} \right) \quad (7.56)$$

Abbreviate $\bigcup_{t_2 \in R_2} \{t_1 \circ t_2 \mid t_1 \in R_1, p(t_1 \circ t_2)\}$ by Y.

Applying the definition of $\Gamma_{G;A}$ yields:

$$\begin{aligned} \{t \circ a \mid t \in \Pi_G(Y), a = (A_1 : e_1, \dots, A_n : e_n), \\ a_i = \text{agg}_i(\{e_i(s) \mid s \in Y, S|_G = t\})\} \end{aligned} \quad (7.57)$$

Compare (20.5) and (20.7). Since $\Pi_G(X) = \Pi_G(Y)$, they can only differ in their values of A_i .

Hence, it suffices to prove that $a_i^2 = a_i$ for $1 \leq i \leq n$ for any given t .

$$\begin{aligned}
a_i^2 &= \text{agg}_i^2(\{s.A_i^1 | s \in X, S|_G = t\}) \\
&= \text{agg}_i^2(\{s.A_i^1 | s \in \bigcup_{t_2 \in R_2} \Gamma_{G_1; A_1}(R_1 \bowtie_p \{t_2\}), S|_G = t\}) \\
&= \text{agg}_i^2(\{s.A_i^1 | s \in \bigcup_{t_2 \in R_2} \{t_1 \circ t_2 \circ a_1 | t_1 \in \Pi_{G_1}(R_1), p(t_1 \circ t_2), \\
&\quad a_1 = (A_1^1 : a_1^1, \dots, A_n^1 : a_n^1) \\
&\quad a_i^1 = \text{agg}_i^1(\{e_i(s_1 \circ t_2) | s_1 \in R_1, S_1|_{G_1=t_1}, p(s_1, t_2)\}) \\
&\quad S|_G = t\}) \\
&= \text{agg}_i^2(\bigcup_{t_2 \in R_2} \{\text{agg}_i^1(\{e_i(s_1 \circ t_2) | t_1 \in \Pi_{G_1}(R_1), p(t_1 \circ t_2), s_1 \in R_1, S_1|_{G_1} = t_1, \\
&\quad p(s_1, t_1), t_1 \circ t_2|_G = t\})\}) \\
&= \text{agg}_i^2(\bigcup_{t_2 \in R_2} \{\text{agg}_i^1(\{e_i(s_1 \circ t_2) | s_1 \in R_1, s_1 \circ t_2|_G = t, p(s_1 \circ t_2)\})\}) \\
&= \text{agg}_i^2(\bigcup_{t_2 \in R_2} \{\text{agg}_i^1(\{e_i(t_1 \circ t_2) | t_1 \in R_1, t_1 \circ t_2|_G = t, p(t_1 \circ t_2)\})\}) \\
&= \text{agg}_i(\bigcup_{t_2 \in R_2} \{e_i(t_1 \circ t_2) | t_1 \in R_1, p(t_1 \circ t_2), t_1 \circ t_2|_G = t\}) \\
&= \text{agg}_i(\{e_i(s) | s \in \bigcup_{t_2 \in R_2} \{t_1 \circ t_2 | t_1 \in R_1, p(t_1 \circ t_2)\}, S|_G = t\}) \\
&= a_i
\end{aligned}$$

Equivalence 20.1 can be used to add an additional coalescing grouping in front of any of a sequence of joins. Consider the schematic operator tree in Figure 20.2(a). It is equivalent to the one in (b), which in turn is equivalent to the one in (c) if the preconditions of Equivalence 20.1 hold. Performing a similar operation multiple times, any of the join operations can be made to be preceded by a coalescing grouping.

7.2.13 ToDo

[611]

7.3 Logical Algebra for Sequences

7.3.1 Introduction

The algebra (NAL) we use here extends the SAL-Algebra [63] developed by Beeri and Tzaban. SAL is the order-preserving counterpart of the algebra used in [161, 162] and in this book.

SAL and NAL work on sequences of sets of variable bindings, i.e., sequences of unordered tuples where every attribute corresponds to a variable. We allow nested tuples, i.e. the value of an attribute may be a sequence of tuples. Single tuples are constructed by using the standard $[\cdot]$ brackets. The concatenation




Figure 7.5: Applications of coalescing grouping

of tuples and functions is denoted by \circ . The set of attributes defined for an expression e is defined as $\mathcal{A}(e)$. The set of free variables of an expression e is defined as $\mathcal{F}(e)$.

The projection of a tuple on a set of attributes A is denoted by $|_A$. For an expression e_1 possibly containing free variables, and a tuple e_2 , we denote by $e_1(e_2)$ the result of evaluating e_1 where bindings of free variables are taken from variable bindings provided by e_2 . Of course this requires $\mathcal{F}(e_1) \subseteq \mathcal{A}(e_2)$. For a set of attributes we define the tuple constructor \perp_A such that it returns a tuple with attributes in A initialized to NULL.

For sequences e we use $\alpha(e)$ to denote the first element of a sequence. We identify single element sequences and elements. The function τ retrieves the tail of a sequence and \oplus concatenates two sequences. We denote the empty sequence by ϵ . As a first application, we construct from a sequence of non-tuple values e a sequence of tuples denoted by $e[a]$. It is empty if e is empty. Otherwise $e[a] = [a : \alpha(e)] \oplus \tau(e)[a]$.

By *id* we denote the identity function. In order to avoid special cases during the translation of XQuery into the algebra, we use the special algebraic operator (\square) that returns a singleton sequence consisting of the empty tuple, i.e., a tuple with no attributes.

We will only define order-preserving algebraic operators. For the unordered counterparts see [162]. Typically, when translating a more complex XQuery into our algebra, a mixture of order-preserving and not order-preserving operators

will occur. In order to keep the paper readable, we only employ the order-preserving operators and use the same notation for them that has been used in [161, 162] and SAL [63].

Again, our algebra will allow nesting of algebraic expressions. For example, within a selection predicate of a select operator we allow the occurrence of further nested algebraic expressions. Hence, a join within a selection predicate is possible. This simplifies the translation procedure of nested XQuery expressions into the algebra. However, nested algebraic expressions force a nested loop evaluation strategy. Thus, the goal of the paper will be to remove nested algebraic expressions. As a result, we perform unnesting of nested queries not at the source level but at the algebraic level. This approach is more versatile and less error-prone.

7.3.2 Algebraic Operators

We define the algebraic operators recursively on their input sequences. For unary operators, if the input sequence is empty, the output sequence is also empty. For binary operators, the output sequence is empty whenever the left operand represents an empty sequence.

The order-preserving selection operator is defined as

$$\hat{\sigma}_p(e) := \begin{cases} \epsilon & \text{if } e = \epsilon \\ \alpha(e) \oplus \hat{\sigma}_p(\tau(e)) & \text{if } p(\alpha(e)) \\ \hat{\sigma}_p(\tau(e)) & \text{else} \end{cases}$$

For a list of attribute names A we define the projection operator as

$$\hat{\Pi}_A(e) := \begin{cases} \epsilon & \text{if } e = \epsilon \\ \alpha(e)|_A \oplus \hat{\Pi}_A(\tau(e)) & \text{else} \end{cases}$$

We also define a duplicate-eliminating projection $\hat{\Pi}_A^D$. Besides the projection, it has similar semantics as the **distinct-values** function of XQuery: it does not preserve order. However, we require it to be deterministic and idempotent. Sometimes we just want to eliminate some attributes. When we want to eliminate the set of attributes A , we denote this by $\hat{\Pi}_{\bar{A}}$. We use $\hat{\Pi}$ also for renaming attributes. Then we write $\hat{\Pi}_{A':A}$. The attributes in A are renamed to those in A' . Attributes other than those in A remain untouched.

The map operator is defined as follows:

$$\hat{\chi}_{a:e_2}(e_1) := \begin{cases} \epsilon & \text{if } e_1 = \epsilon \\ \alpha(e_1) \circ [a : e_2(\alpha(e_1))] \oplus \hat{\chi}_{a:e_2}(\tau(e_1)) & \text{else} \end{cases}$$

It extends a given input tuple $t_1 \in e_1$ by a new attribute a whose value is computed by evaluating $e_2(t_1)$. For an example see Figure ??.

We define the cross product of two tuple sequences as

$$e_1 \hat{\times} e_2 := \begin{cases} \epsilon & \text{if } e_1 = \epsilon \\ (\alpha(e_1) \hat{\times} e_2) \oplus (\tau(e_1) \hat{\times} e_2) & \text{else} \end{cases}$$

	R_2			$\hat{\chi}_{a:\hat{\sigma}_{A_1=A_2}(R_2)}(R_1) =$
R_1	A_2	B	A_1	a
A_1	1	2	1	$\langle [1, 2], [1, 3] \rangle$
1	1	3	2	$\langle [2, 4], [2, 5] \rangle$
2	2	4	3	$\langle \rangle$
3	2	5		

Figure 7.6: Example for Map Operator

where

$$e_1 \hat{\times} e_2 := \begin{cases} \epsilon & \text{if } e_2 = \epsilon \\ (e_1 \circ \alpha(e_2)) \oplus (e_1 \hat{\times} \tau(e_2)) & \text{else} \end{cases}$$

We are now prepared to define the join operation on ordered sequences:

$$e_1 \bowtie_p e_2 := \sigma_p(e_1 \hat{\times} e_2)$$

We define the semijoin as

$$e_1 \hat{\bowtie}_p e_2 := \begin{cases} \alpha(e_1) \oplus (\tau(e_1) \hat{\bowtie}_p e_2) & \text{if } \exists x \in e_2 \ p(\alpha(e_1) \circ x) \\ \tau(e_1) \hat{\bowtie}_p e_2 & \text{else} \end{cases}$$

and the anti-join as

$$e_1 \hat{\delta}_p e_2 := \begin{cases} \alpha(e_1) \oplus (\tau(e_1) \hat{\delta}_p e_2) & \text{if } \nexists x \in e_2 \ p(\alpha(e_1) \circ x) \\ \tau(e_1) \hat{\delta}_p e_2 & \text{else} \end{cases}$$

The left outer join, which will play an essential role in unnesting, is defined as $e_1 \hat{\bowtie}_p^{g:e} e_2 :=$

$$\begin{cases} (\alpha(e_1) \bowtie_p e_2) \oplus (\tau(e_1) \hat{\bowtie}_p^{g:e} e_2) & \text{if } (\alpha(e_1) \bowtie_p e_2) \neq \epsilon \\ (\alpha(e_1) \circ \perp_{\mathcal{A}(e_2)} \{g\}) \circ [g : e] & \text{else} \\ \oplus (\tau(e_1) \hat{\bowtie}_p e_2) & \end{cases}$$

where $g \in \mathcal{A}(e_2)$. Our definition deviates slightly from the standard left outer join operator, as we want to use it in conjunction with grouping and (aggregate) functions. Consider the relations R_1 and R_2 in Figure ???. If we want to join R_1 (via left outer join) to R_2^{count} that is grouped by A_2 with counted values for B , we need to be able to handle empty groups (for $A_1 = 3$). e defines the value given to attribute g for values in e_1 that do not find a join partner in e_2 (in this case 0).

We define the dependency join (d-join for short) as

$$e_1 \hat{\triangleleft} e_2 \hat{\triangleright} := \begin{cases} \epsilon & \text{if } e_1 = \epsilon \\ \alpha(e_1) \hat{\times} e_2(e_1) \oplus \tau(e_1) \hat{\triangleleft} e_2 \hat{\triangleright} & \text{else} \end{cases}$$

Let $\theta \in \{=, \leq, \geq, <, >, \neq\}$ be a comparison operator on atomic values. The grouping operator which produces a sequence-valued new attribute containing “the group” is defined by using a binary grouping operator.

$$\hat{\Gamma}_{g;\theta A;f}(e) := \hat{\Pi}_{A:A'}(\hat{\Pi}_{A':A}^D(\hat{\Pi}_A(e))\hat{\Gamma}_{g;A'\theta A;f}e)$$

where the binary grouping operator (sometimes called nest-join [720]) is defined as

$$e_1 \hat{\Gamma}_{g;A_1\theta A_2;f} e_2 := \begin{cases} \epsilon & \text{if } e_1 = \epsilon \\ \alpha(e_1) \circ [g : G(\alpha(e_1))] \oplus (\tau(e_1) \hat{\Gamma}_{g;A_1\theta A_2;f} e_2) & \text{else} \end{cases}$$

Here, $G(x) := f(\sigma_{x|_{A_1}\theta A_2}(e_2))$ and function f assigns a meaningful value to empty groups. See also Figure ?? for an example. The unary grouping operator processes a single relation and obviously groups only on those values that are present. The binary grouping operator works on two relations and uses the left hand one to determine the groups. This will become important for the correctness of the unnesting procedure.

R_1	R_2	$\hat{\Gamma}_{g;=A_2;count}(R_2) =$	$\hat{\Gamma}_{g;=A_2;id}(R_2) =$
$\begin{array}{ c } \hline A_1 \\ \hline 1 \\ 2 \\ 3 \\ \hline \end{array}$	$\begin{array}{ c c } \hline A_2 & B \\ \hline 1 & 2 \\ 1 & 3 \\ 2 & 4 \\ 2 & 5 \\ \hline \end{array}$	$\begin{array}{ c c } \hline A_2 & g \\ \hline 1 & 2 \\ 2 & 2 \\ \hline \end{array}$	$\begin{array}{ c c } \hline A_2 & g \\ \hline 1 & \langle [1, 2], [1, 3] \rangle \\ 2 & \langle [2, 4], [2, 5] \rangle \\ \hline \end{array}$

$R_1 \hat{\Gamma}_{g;A_1=A_2;id}(R_2) =$
$\begin{array}{ c c } \hline A_1 & g \\ \hline 1 & \langle [1, 2], [1, 3] \rangle \\ 2 & \langle [2, 4], [2, 5] \rangle \\ 3 & \langle \rangle \\ \hline \end{array}$

Figure 7.7: Examples for Unary and Binary $\hat{\Gamma}$

Given a tuple with a sequence-valued attribute, we can unnest it using the unnest operator defined as

$$\hat{\mu}_g(e) := \begin{cases} \epsilon & \text{if } e = \epsilon \\ (\alpha(e)|_{\overline{\{g\}}} \hat{\times} \alpha(e).g) \oplus \hat{\mu}_g(\tau(e)) & \text{else} \end{cases}$$

where $e.g$ retrieves the sequence of tuples of attribute g . In case that g is empty, it returns the tuple $\perp_{\mathcal{A}(e.g)}$. (In our example in Figure ??, $\hat{\mu}_g(R_2^g) = R_2$.)

We define the unnest map operator as follows:

$$\hat{\Upsilon}_{a:e_2}(e_1) := \hat{\mu}_g(\chi_{g:e_2[a]}(e_1))$$

This operator is mainly used for evaluating XPath expressions. Since this is a very complex issue [285, 287, 357], we do not delve into optimizing XPath evaluation but instead take an XPath expression occurring in a query as it is and use it in the place of e_2 . Optimized translation of XPath is orthogonal to our unnesting approach and not covered in this paper. The interested reader is referred to [357].

7.3.3 Equivalences

To acquaint the reader with ordered sequences, we state some familiar equivalences that still hold.

$$\hat{\sigma}_{p_1}(\hat{\sigma}_{p_2}(e)) = \hat{\sigma}_{p_2}(\hat{\sigma}_{p_1}(e)) \quad (7.58)$$

$$\hat{\sigma}_p(e_1 \hat{\times} e_2) = \hat{\sigma}_p(e_1) \hat{\times} e_2 \quad (7.59)$$

$$\hat{\sigma}_p(e_1 \hat{\times} e_2) = e_1 \hat{\times} \hat{\sigma}_p(e_2) \quad (7.60)$$

$$\hat{\sigma}_{p_1}(e_1 \bowtie_{p_2} e_2) = \hat{\sigma}_{p_1}(e_1) \bowtie_{p_2} e_2 \quad (7.61)$$

$$\hat{\sigma}_{p_1}(e_1 \bowtie_{p_2} e_2) = e_1 \bowtie_{p_2} \hat{\sigma}_{p_1}(e_2) \quad (7.62)$$

$$\hat{\sigma}_{p_1}(e_1 \hat{\times}_{p_2} e_2) = \hat{\sigma}_{p_1}(e_1) \hat{\times}_{p_2} e_2 \quad (7.63)$$

$$\hat{\sigma}_{p_1}(e_1 \hat{\bowtie}_{p_2}^{g:e} e_2) = \hat{\sigma}_{p_1}(e_1) \hat{\bowtie}_{p_2}^{g:e} e_2 \quad (7.64)$$

$$e_1 \hat{\times} (e_2 \hat{\times} e_3) = (e_1 \hat{\times} e_2) \hat{\times} e_3 \quad (7.65)$$

$$e_1 \bowtie_{p_1}(e_2 \bowtie_{p_2} e_3) = (e_1 \bowtie_{p_1} e_2) \bowtie_{p_2} e_3 \quad (7.66)$$

$$\hat{\sigma}_p(e_1 \hat{\times} e_2) = e_1 \bowtie_p e_2 \quad (7.67)$$

$$e_1 \hat{<} e_2 \hat{>} = e_1 \hat{\times} e_2 \quad (7.68)$$

$$\hat{\Upsilon}_{a:f(\chi_b(e))}(\square) = \hat{\Pi}_{a:b}(f(e)) \quad (7.69)$$

$$\hat{\Upsilon}_{a:e_2}(e_1) = e_1 \hat{\times} \hat{\Upsilon}_{a:e_2}(\square) \quad (7.70)$$

$$(7.71)$$

Of course, in the above equivalences the usual restrictions hold. For example, if we want to push a selection predicate into the left part of a join, it may not reference attributes of the join's right argument. In other words, $\mathcal{F}(p_1) \cap \mathcal{A}(e_2) = \emptyset$ is required. As another example, equivalence 7.70 only holds if $\mathcal{F}(e_1) \cap \mathcal{A}(e_1) = \emptyset$. In Eqv. 7.69 the function f may not alter the schema and b must be an attribute name. Please note that cross product and join are still associative in the ordered context. However, neither of them is commutative. Further, pushing selections into the second argument of a left-outer join is (in general) not possible. For strict predicates we can do better but this is beyond the scope of the paper.

7.3.4 Bibliography

Zaniolo [458]

7.4 Literature

- NF²: [4, 174, 371, 636, 637, 469, 638, 470, 666]
- HAS: [107]
- Aggregates: [422, 431]
- SQL to Algebra: [783, 110]
- Calculus to Algebra: [783, 542]

- BAGs: [22]
- Algebra with control over duplicate elimination: [189]
- OO Algebra of Steenhagen et al. [720, 721, 719, 722]
- OO Algebra of Cluet and Moerkotte [161, 162].
- OO Algebra [147]
- OO Algebra [160]
- OO Algebra [325]
- OO Algebra [484]
- OO Algebra [650]
- OO Algebra [693, 692, 691, 694]
- OO Algebra [779, 780]
- OO Algebra [841]
- SAL [63]: works on lists. Intended for semistructured data. SAL can be thought of as the order-preserving counterpart of the algebra presented in [161, 162] extended to handle semistructured data. These extensions are similar to those proposed in [5, 153]
- TAX [393]: The underlying algebra's data model is based on sets of ordered labeled trees. Intended for XML.
- XML: Construction: [242, 243]
- no-name [711]: order preserving algebra for OLAP-style queries
- [316]
- Document Processing Algebras: [154, 329]
- Geo: [330]

Chapter 8

Declarative Query Representation

8.1 Calculus Representations

relational calculus originally introduced by Codd: [168, 167].

Variant for embedding in Pascal/R: [665]

calculus for complex objects: [44]

8.2 Datalog

8.3 Tableaux Representation

Tableaus have been introduced by [18, 19, 20] Tableaus are able to represent a particular kind of queries, the so called *conjunctive queries* ([118], [623]).

Expressions containing disjunction (set union) and negation (set difference) can be represented by sets of tableaus ([649],[403]).

query processing with tables: [559]

8.4 Monoid Comprehension

[96, 238, 239]

8.5 Expressiveness

transitivity: [588]. aggregates: [431]. complex object and nested relations: [3].

8.6 Bibliography

Chapter 9

Translation and Lifting

9.1 Query Language to Calculus

9.2 Query Language to Algebra

9.3 Calculus to Algebra

9.4 Algebra to Calculus

9.5 Bibliography

Chapter 10

Query Equivalence, Containment, Minimization, and Factorization

This chapter is devoted to three related problems. Let q be a query and d be a database. We denote by $q(d)$ the result of evaluating the query on the database. Two queries q_1 and q_2 are by definition *equivalent* ($q_1 \equiv q_2$), if $q_1(d) = q_2(d)$ for all databases d . The problem of deciding whether two queries are equivalent is the *query equivalence* problem. We define $q_1 \subseteq q_2$ if $q_1(d) \subseteq q_2(d)$ for all databases d . The problem of deciding whether for two $q_1 \subseteq q_2$ holds is the *query containment* problem. Obviously, $q_1 \equiv q_2 \iff q_1 \subseteq q_2 \wedge q_2 \subseteq q_1$.

The third problem is query minimization. Consider the following conjunctive query written as a rule in datalog:

$$q(X, Y) : -p(X, Y), p(X, Y).$$

Obviously, this query is equivalent to

$$q(X, Y) : -p(X, Y).$$

under set semantics. The latter query now contains fewer body literals. Query minimization now asks for an equivalent query with the least possible number of body literals. One possible approach is to successively delete a body literal until no more body literal can be deleted without violating equivalence to the original query.

The above example is also illustrative since it shows that query equivalence (and thus query containment) differs under different semantics: whereas the above two queries are equivalent under set semantics, they are not under bag semantics. To see this, consider the extensional database $\{p(a, b), p(a, b)\}$. The result of the first query contains $p(a, b)$ four times whereas the last query contains is only 2 times.

10.1 Set Semantics

10.1.1 Conjunctive Queries

A conjunctive query may contain literals where the predicate is built-in. For example, it applies comparison operators $=, <, >, \leq, \geq, \neq$. Any explicit use of equality can be eliminated as follows. For any literal of the form $X = c$, any occurrence of X is replaced by c and the equality literal is dropped from the query clause. For any literal of the form $X = Y$, any occurrence of Y is replaced by X and the equality literal is dropped from the query clause. This procedure is not possible for the other comparison operators $<, >, \leq, \geq, \neq$ which we call *inequality operators*. An *inequality* is any literal using an inequality operator.

Containment and minimization for conjunctive queries without inequalities are NP-complete Problems. First note that a tableau directly corresponds to a conjunctive query with all body literals having a common predicate. From that and the NP-completeness results for tableau containment which in turn follows immediately from first order subsumption [58, 272], it follows immediately that containment of conjunctive queries is NP-complete. Chandra and Merlin proved that minimization is NP-complete [118]. The complexity of checking for equivalence of conjunctive queries is related to graph isomorphism.

EX

The procedure for checking query containment builds upon mappings from queries to queries¹. These mappings have two different names: *homomorphism* and *containment mapping*.

Let q_1 and q_2 be the two queries for which we want to check containment. Assume the q_i are of the form

$$\begin{aligned} q_1 : r_1 & :- l_1, \dots, l_k \\ q_2 : r_2 & :- l'_1, \dots, l'_m \end{aligned}$$

Let $\mathcal{V}(q_i)$ be the set of variables occurring in q_i , and $\mathcal{C}(q_i)$ be the set of constants occurring in q_i . Further, let h be a substitution $h : \mathcal{V}(q_2) \rightarrow (\mathcal{V}(q_1) \cup \mathcal{C}(q_1))$. We call h a *containment mapping* from q_2 to q_1 , if and only if the following conditions are fulfilled:

1. $h(r_2) = r_1$ for the head literals, and
2. for all i ($1 \leq i \leq m$) there exists a j ($1 \leq j \leq k$) such that $h(l'_i) = l_j$.

The latter condition states that for each body literal l'_i in q_2 there is a body literal l_j in q_1 such that $h(l'_i) = l_j$. Note that this does not imply that h is injective or surjective.

The following theorem connects containment mappings with the containment problem:

Theorem 10.1.1 *Let q_1 and q_2 be two conjunctive queries. Then $q_1 \subseteq q_2$ if and only if there is a containment mapping h mapping q_2 to q_1 .*

¹In fact, Chandra and Merlin mapped *natural models* which is essentially the same.

Consider the following example:

$$\begin{aligned} q_1 : p(X_1, X_2) & : - q(X_2, X_1), q(X_1, X_3) \\ q_2 : p(Y_1, Y_2) & : - q(Y_2, Y_1), q(Y_3, Y_1), q(Y_1, Y_4) \end{aligned}$$

Consider h with $h(Y_1) = X_1$, $h(Y_2) = X_2$, $h(Y_3) = X_2$, and $h(Y_4) = X_3$. Then

$$\begin{array}{l} l : p(Y_1, Y_2) \quad \left| \quad \begin{array}{ccc} q(Y_2, Y_1) & q(Y_3, Y_1) & q(Y_1, Y_4) \\ q(X_2, X_1) & q(X_2, X_1) & q(X_1, X_3) \end{array} \\ h(l) : p(X_1, X_2) \end{array}$$

and, hence, $q_1 \subseteq q_2$.

A query q is minimal, if it contains the minimal possible number of body literals. More formally, q is *minimal*, if for any query q' with $q \equiv q'$ the number of body literals in q' greater than or equal to the number of body literals in q . The following theorem shows that our initial thoughts on minimization are correct for conjunctive queries.

Theorem 10.1.2 *Let q be a conjunctive query. Then there is a minimal query q' equivalent to q such that q' results from q by deleting zero or more body literals.*

This suggests a simple procedure for minimizing a given query q . For every body literal check whether some containment mapping h exists such that it is subsumed by some other body literal. Note that this containment mapping must not rename head variables.

Let q and q' be two conjunctive queries. If q can be derived from q' solely by reordering body literals and renaming variables, then q and q' are called *isomorphic*. Minimal queries are unique up to some isomorphism. Obviously, minimizing conjunctive queries is also NP-complete.

Let us now come to unions of conjunctive queries. Let $Q = Q_1 \cup \dots \cup Q_k$ and $Q' = Q'_1 \cup \dots \cup Q'_l$ be two unions of conjunctive queries Q_i and Q'_j with a common head predicate. A containment mapping h from Q to Q' maps each Q_i to some Q'_j such that $h(Q_i) \subseteq Q'_j$. Sagiv and Yannakakis showed the following theorem [649].

Theorem 10.1.3 *Let $Q = Q_1 \cup \dots \cup Q_k$ and $Q' = Q'_1 \cup \dots \cup Q'_l$ be two unions of conjunctive queries Q_i and Q'_j with a common head predicate. Then $Q \subseteq Q'$ if and only if there is a containment mapping from Q to Q' .*

This theorem gives us a corollary which allows us minimizing unions of conjunctive queries by a pairwise checking of containment [649] (see also [767]).

Corollary 10.1.4 *Let $Q = Q_1 \cup \dots \cup Q_k$ be a union of conjunctive queries Q_i with common head predicate. Then there exists a subset R of Q such that*

1. $R \equiv Q$
2. $\neg \exists R' \subset R \quad R' \equiv Q$
3. If Q_m is any equivalent to Q , then there is a containment mapping from Q_m to R but none from Q_m to any proper subset R' of R .

This corollary implies that we can minimize a query that is a union of conjunctive queries by eliminating those conjunctive queries Q_i from it that are contained in some Q_j .

For conjunctive queries the problems of containment, equivalence, and minimization are

The problems of containment, equivalence, and minimization of conjunctive queries are most difficult if all body literals have a common predicate p . This is quite an unrealistic assumption as typical conjunctive queries will not only self-join the same relation. A first question is thus whether there exist special cases where there are polynomial algorithms for containment checking. Another strain of work is devoted to more complex queries. As it turns out, the results become less nice and more restricted.

10.1.2 ... with Inequalities

We now turn to conjunctive queries with inequalities in their body. For this section, we assume that the domain is totally ordered and dense. That is, for all x and y with $x < y$, there exists a z with $x < z < y$. In this context, we have the following theorem:

Theorem 10.1.5 *Assume the two conjunctive queries q_1 and q_2 are of the form*

$$\begin{aligned} q_1 : p_1 & : - l_1, \dots, l_k, e_1, \dots, e_l \\ q_2 : p_2 & : - l'_1, \dots, l'_m, e'_1, \dots, e'_n \end{aligned}$$

where p_i are the head literals, l_i and l'_i are ordinary subgoals and e_i and e'_i are inequalities. Let h be a containment mapping from q_2 to q_1 where both are restricted to their ordinary literals. If additionally for all $i = 1, \dots, n$ we have

$$e_1, \dots, e_l \implies h(e'_i)$$

then $q_1 \subseteq q_2$.

This result is due to Klug [432] who used the following procedure to reason about inequalities using comparison operators in $\{=, <, \leq\}$. Given a set of inequalities L , an directed graph G is defined whose nodes are the variables and constants in L . Whenever for all $x < y$ or $x \leq y$ in L , the edge (x, y) is added to G . For all constants c and c' in L , if $c < c'$ then we add an edge (c, c') . Edges are labeled with the according comparison operator. For equality predicates, an edge in both direction is added. Given the graph G , we conclude that $x \leq y$ if there is a path from x to y and $x < y$ only if additionally at least one edge is labelled by $<$. An alternative is to use the procedure presented in Section 11.2.3 to solve the inequality inference problem. It also allows for the comparison operator \neq .

To see why a dense domain is important consider the domain of integers. From $1 < x < 3$ we can easily conclude that $x = 2$, a fact we can derive neither from the procedure above nor from the axioms and inference procedure presented in Section 11.2.3.

Unfortunately, the other direction of Theorem ?? is wrong as the following example shows:

$$\begin{aligned} q_1 : p(X_1, X_2) & : - q(X_1, X_2), r(X_3, X_4), r(X_4, X_3) \\ q_2 : p(Y_1, Y_2) & : - q(Y_1, Y_2), r(Y_3, Y_4), Y_3 \leq Y_4 \end{aligned}$$

Obviously, $Y_3 \leq Y_4$ cannot be implied by any non-existing inequalities from q_1 . However, for q_1 to be non-empty, we must have $r(a, b)$ and $r(b, a)$ for some a and b . We also have $a \leq b$ or $b \leq a$. In the former case, we can chose $Y_3 = a$ and $Y_4 = b$ and in the latter $Y_3 = b$ and $Y_4 = a$ to satisfy $r(Y_3, Y_4)$ and $Y_3 \leq Y_4$.

Klug provides an alternative method to solve the containment problem. It builds upon *canonical models*. He then shows that if the containment test succeeds for all canonical models then and only then query containment holds [432]. Klug does not give an explicit algorithm for constructing these canonical models but these can be found in the literature []. He also gives two simple subclasses of inequality queries, where constants are allowed only on the left- or only on the right-hand side. For these subclasses if and only if holds in the above Theorem ??.

Although the theorem is stated in terms of conjunctive queries with inequalities, it holds for any predicate p . Assume two queries of the following form:

$$\begin{aligned} q_1 : p_1 & : - l_1, \dots, l_k, P \\ q_2 : p_2 & : - l'_1, \dots, l'_m, P' \end{aligned}$$

where P and P' are arbitrary formulas. If there is a containment mapping from q_2 to q_1 where both are restricted to their ordinary literals and $P \implies h(P')$, then $q_1 \subseteq q_2$.

10.1.3 ... with Negation

The first incarnation of negation we consider is set difference. Here, Sagiv and Yannakakis where the first to derive some results [649].

[171]

10.1.4 ... under Constraints

constraints: [401, 402, 378]

negation+constraints: [232, 233, 234, 235]

10.1.5 ... with Aggregation

[172, 173]

10.2 Bag Semantics

10.2.1 Conjunctive Queries

- definition bag-containment, bag-equivalence [189, 427]
 - characterizations [135, 383] (no proofs in [136])

- complexity results [135]
- definition bag-set containment, bag-set equivalence [135]

10.3 Sequences

10.3.1 Path Expressions

XPath constructs and their short-hands to denote XPath sublanguages.

- branching ('[]')
- wild cards ('*')
- descendant axis ('//')
- disjunction : only binary *or* branching, ('—') (or-branching)

Otherwise XPath only contains the child axis and node name tests. These sublanguages are represented as tree patterns.

Query containment for certain subclasses:

- $XP^{\square,*,//}$ is coNP-complete [518]

Consider we have to answer $p \subseteq p'$. Then

- if $p \in P^{\square,//}$ and $p' \in P^{\square,*,//}$ then query containment is coNP-complete.
- in PTIME if number of '/' is restricted by d which then gives the degree of the polynomial describing the time complexity
- remains coNP-complete if p contains no '*' p' contains at most two '*'_s
- remains coNP-complete if p contains at most 5 branches and p' contains at most 3 branches.
- $P^{\square,*}$ is in PTIME (follows from work on conjunctive acyclic queries [830], also noted by Wood [811])
- $P^{\square,//}$ is in PTIME [27]
- $P^{*,//}$ is in PTIME (these are related to a fragment of regular expressions [519])
- P^{or} is in PTIME
- $P^{\square,or}$ is coNP-complete
- P^{\square} is coNP-complete [518]
[545] showed that $P^{\square,*,//,\square}$ is coNP-complete for infinite alphabets and in PSPACE for finite alphabets.

- $P^{//,|}$ is PSPACE-complete
- $P^{\square,*,//}$ with variable binding and equality tests is Π_2^P -hard [201]

A PTIME algorithm for the fragment $P^{//}$ can be found in [95]

Florescu, Levy, and Suciu showed that for a language quite similar to $P^{\square,//}$ containment is NP-complete if evaluated on a graph-based data model instead of a tree-based one [246].

Calvanese et al. also consider a graph-based data model and more expressive queries [99].

[545] also contains work on languages with variable bindings with different semantics.

More result: [201]

query containment with DTDs: [812]

Schwentick gives a very good overview over complexity results for containment checking [667]. We should repeat his table here.

10.4 Minimization

minimization: [445]

10.5 Detecting common subexpressions

[244, 337, 335]

10.5.1 Simple Expressions

Simple Non-Expensive Expressions

Simple Expensive Expressions

10.5.2 Algebraic Expressions

10.6 Bibliography

In a pair of papers Aho, Sagiv, and Ullmann [16, 17] study equivalence, containment, and minimization problems for tableaux. More specifically, they introduce a restricted variant of relational expressions containing projection, natural join, and selection with predicates that only compare attributes with constants. They further assume the existence of a universal relation. That is, every relation R is the projection of the universal relation on $\mathcal{A}(R)$. Now, these restricted conjunctive queries can be expressed with tableaux. The authors tableaux equivalence, containment, and minimization problems also in the presence of functional dependences. The investigated problems are all NP-complete. Since the practical usefulness is limited we do not give the concrete results of this pair of papers.

[134, 137] contains (complexity) results for deciding query equivalence in the case of recursive and nonrecursive datalog.

View selection problem (just pointers): ...

[398, 395, 396]

[141]

- conjunctive queries: equivalence and minimization are NP-complete [118, 17]
In [17] tableaux are used.
- polynomial algorithms for equivalence and minimization for simple tableaux: [17, 16]
- union of elementary differences: Π_2^P complete: remark in [649] and a pointer to the thesis of Sagiv [646].
- acyclic conjunctive queries: PTIME [830]
- equivalence $(\sigma, \bowtie, \pi, \cup)$, equivalence $(\sigma, \bowtie, \pi, \cup, \setminus)$: Π_2^P -complete [649]
- Recursive Datalog: [85]

Part III

Rewrite Techniques

Chapter 11

Simple Rewrites

11.1 Simple Adjustments

11.1.1 Rewriting Simple Expressions

Constant Folding

Constant subexpressions are evaluated and the result replaces the subexpression. For example an expression $1/100$ is replaced by 0.01 . Other expressions like $a - 10 = 50$ can be rewritten to $a = 40$. However, the latter kind of rewrite is rarely performed by commercial systems.

Eliminate BETWEEN

A predicate of the form Y BETWEEN X AND Z is replaced by $X \leq Y$ AND $Y \leq Z$. This step not only eliminates syntactic sugar but also enables transitivity reasoning to derive new predicates (see).

Eliminate IN

A predicate of the form x IN (c_1, \dots, c_n) is rewritten to $x = c_1$ OR \dots OR $x = c_n$. This eliminates on form of the IN predicate and enables multikey index access.

Another possibility is to use a table function that produces a table with one column whose values are exactly those in the IN-list. From thereon, regular optimization takes place. This possibility is also investigated when several comparisons of a column with a constants are disjunctively connected.

Eliminating LIKE

A predicate of the form a LIKE 'Guy' can only be rewritten to $a = 'Guy'$ if a is of type varchar. This is due to the different white space padding rules for LIKE and =.

Start and Stop conditions derived from LIKE predicates

A predicate of the form a LIKE 'bla%' gives rise to a start condition $a \geq$ 'bla'. Which can enable subsequent index usage. A stop predicate of the form $a <$ 'blb' can also be derived. completing a range predicate for an index scan. Start and stop conditions can only be derived if there is no leading '%' in the pattern.

Pushing NOT operations down and eliminating them

NOT operations need to be pushed downwards for correctness reasons. Attention has to be paid to the IS NOT NULL and IS NULL predicates. XXX complete set of rules go into some table.

Merge AND, OR, and other associative operations

While parsing, AND and OR operations are binary. For simpler processing they are often n-ary in the internal representation. Therefor $(p \text{ AND } (q \text{ AND } r))$ is rewritten to $(\text{AND } p \ q \ r)$.

In general, associative nested operations should be merged. Examples of other associative operations are $+$ and $*$.

Normalized Argument Order for Commutative Operations

ToDo

enabling factorization, constant folding: move constants to the left Speed up evaluation of *equal*.

Eliminate - and /

$$(x - y) \rightsquigarrow x + (-y) \quad x/y \rightsquigarrow x * (1/y)$$

Adjust join predicates

$A = B + C$ becomes $A - C = B$ if A and B are from one relation and C is from another.

Simplifying boolean expressions

The usual simplification rules for boolean expressions can be applied. For example, if a contradiction can be derived.

Eliminating ANY, SOME, and ALL

ANY and SOME operators in conjunction with a comparison operator are rewritten into disjunction of comparison predicates. For example $a >$ ANY (c_1, c_2) is rewritten to $a > c_1$ OR $a > c_2$. Correspondingly, an ALL operator with a constant list is rewritten into a conjunction of comparisons. For example, $a >$ ALL (c_1, c_2) is rewritten to $a > c_1$ AND $a > c_2$.

If a subquery occurs, then the ANY or SOME expression is rewritten to a correlated subquery in an EXIST predicate. Consider the query $a >$ ANY

(SELECT b FROM ...WHERE p). It is rewritten to EXISTS(SELECT ...FROM ...WHERE p AND a > b).

Correspondingly, a subquery within an ALL operator is rewritten into a NOT EXISTS subquery. For example, a > (SELECT b FROM ...WHERE p) is rewritten into NOT EXISTS (SELECT b FROM ...WHERE p and a <= b)

- CASE $j=i$ UNION

11.1.2 Normal forms for queries with disjunction

Another step of the NFST component or the first step of the rewriting component can be the transformation of boolean expressions found in *where* clauses in order to account for NULL values. Pushing **not** operators inside the boolean expression allows to use *two-valued logic* instead of *three-valued logic*. If we miss this step, we can get wrong results.

Another possible step is the subsequent transformation of the boolean expressions in **where** clauses into disjunctive normal form (DNF) or conjunctive normal form (CNF). This step is not always necessary and really depends on which plan generation approach is taken. Hence, this step could take place as late as in a preparatory step for plan generation. It is (obviously) only necessary if the query contains disjunctions. We discuss plan generation for queries with disjunctions in Section ??.

11.2 Deriving new predicates

Given a set of conjunctive predicates, it is often possible to derive new predicates which might be helpful during query plan generation.

This section discusses ways to infer new predicates.

11.2.1 Collecting conjunctive predicates

A query predicate may not only contain the **and** connector, but also **or** or **not**.

For the inference rules in this section we need base predicates that *occur conjunctively*.

We say that a (base) predicate q occurs conjunctively in a (complex) predicate p if $p [q \leftarrow false]$ can be simplified to *false*. That is, if we replace every occurrence of q by *true*, the simplification rules in Figure 11.1 (Fig. ??) simplify $p [q \leftarrow true]$ to *false*.

These simplification rules can be used to implement a simple member function *occursConjunctively* to determine whether a predicate occurs conjunctively in a predicate or not. Together with a member function or visitor *Collect-BasePredicates*, we can compute the set of conjunctively occurring predicates. This set will form the basis for the next subsections.

11.2.2 Equality

Equality is a reflexive, symmetric and transitive binary relationship (see Fig. 11.2). Such a relation is called an *equivalence relation*. Hence, a set of conjunctively

$$\begin{aligned}
NOT\ true &\rightarrow false \\
NOT\ false &\rightarrow true \\
p\ AND\ true &\rightarrow p \\
p\ AND\ false &\rightarrow false \\
p\ OR\ true &\rightarrow true \\
p\ OR\ false &\rightarrow p
\end{aligned}$$

Figure 11.1: Simplification rules for boolean expressions

$$\begin{aligned}
&x = x \\
x = y &\implies y = x \\
x = y \wedge y = z &\implies x = z
\end{aligned}$$

Figure 11.2: Axioms for equality

occurring equality predicates implicitly partitions the set of composed terms (IUs) into disjunctive equivalence classes.

Constants: Let X be an equivalence class of equal expressions. Let Y be the set of all equality expressions that contributed to X . Then, in the query predicate we replace all expressions $x = y$ by $x = c$ and $y = c$ and subsequently eliminate redundant expressions.

$$\sigma_{x=c}(e_1 \bowtie_{x=y} e_2) \equiv \sigma_{x=c}(e_1) \times \sigma_{y=c}(e_2)$$

replace all predicates by $IU=C.IU$'s equivalent to a constant In [178] an abstract data structure is presented that helps computing the equivalence classes fast and also allows for a fast check whether two terms (IUs) are in the same equivalence class. Since we are often interested in whether a given IU is equal to a constant - or, more specifically, equal to another IU bound to a constant -, we have to modify these algorithms such that the IU bound to a constant, if it exists, becomes the representative of its equivalence class.

For the member functions *addEqualityPredicate*, *getEqualityRepresentative* and *isInSameEqualityClass* we need an attribute *_equalityRepresentative* in class IU that is initialized such that it points to itself. Another member *_equalityClassRank* is initialized to 0. The code for the two member functions is given in Figure 11.3.

By calling *addEqualityPredicate* for all conjunctively occurring equality predicates we can build the equivalence classes.

11.2.3 Inequality

Table 11.1 gives a set of axioms used to derive new predicates from a set of conjunctively occurring inequalities S (see [767], see Fig. 11.4).

These axioms have to be applied until no more predicates can be derived. The following algorithm [767] performs this task efficiently:

1. Convert each $X < Y$ into $X \neq Y$ and $X \leq Y$.
2. Compute the transitive closure of \leq .
3. Apply axiom A8 until no more new predicates can be derived.
4. Reconstruct $<$ by using axiom A4.

Step 3 can be performed as follows. For any true IUs X and Y we find these IUs Z with $X \leq Z \leq Y$.

Then we check whether any two such Z 's are related by \neq . Here, it is sufficient to check the original \neq pairs in S and these derived in 1.

A1	: $X \leq X$
A2	: $X < Y \Rightarrow X \leq Y$
A3	: $X < Y \Rightarrow X \neq Y$
A4	: $X \leq Y \wedge X \neq Y \Rightarrow X < Y$
A5	: $X \neq Y \Rightarrow Y \neq X$
A6	: $X < Y \wedge Y < Z \Rightarrow X < Z$
A7	: $X \leq Y \wedge Y \leq Z \Rightarrow X \leq Z$
A8	: $X \leq Z \wedge Z \leq Y \wedge X \leq W \wedge W \leq Y \wedge W \neq Z \Rightarrow X \neq Y$

Table 11.1: Axioms for inequality

11.2.4 Aggregation

Let R_1, \dots, R_n be relations or views, A_1, \dots, A_m attributes thereof, p_w and p_h predicates, and a_1, \dots, a_l expressions of the form $f_j(B_j)$ for aggregate functions f_j and attributes B_j . For a query block of the form

```

select   $A_1, \dots, A_k, a_1, \dots, a_l$ 
from     $R_1, \dots, R_n$ 
where    $p_w$ 
group by  $A_1, \dots, A_m$ 
having   $p_h$ 

```

we consider the derivation of new predicates [474]. Obviously, the following predicates are true:

$$\begin{aligned} \min(B) &\leq B \\ \max(B) &\geq B \\ \max(B) &\geq \min(B) \\ \min(B) &\leq \text{avg}(B) \\ \text{avg}(B) &\leq \max(B) \end{aligned}$$

If p_w contains conjunctively a predicate $B\theta c$ for some constant c , we can further infer

$\min(B) \theta c$	if $\theta \in \{>, \geq\}$	These predicates can then
$\max(B) \theta c$	if $\theta \in \{<, \leq\}$	
$\text{avg}(B) \theta c$	if $\theta \in \{<, \leq, >, \geq\}$	

be used to derive further predicates. The original and the derived predicates are useful when the query block is embedded in another query block since we are allowed to add them to the embedding query block conjunctively (see Section 12.3).

If we know restrictions on the aggregates from some embedding query block, we might be able to add predicates to p_w . The following table contains the restrictions on an aggregate we know in the left column and the predicates we can infer in the right column:

$\max(B) \geq c$	\rightsquigarrow	$B \geq c$	if no other aggregation occurs
$\max(B) > c$	\rightsquigarrow	$B > c$	if no other aggregation occurs
$\min(B) \leq c$	\rightsquigarrow	$B \leq c$	if no other aggregation occurs
$\min(B) < c$	\rightsquigarrow	$B < c$	if no other aggregation occurs

Note that the aggregation occurring in the left column must be the only aggregation found in the query block. That is, $l = 1$ and p_h contains no aggregation other than a_1 . To see why this is necessary, consider the following query

```
select deptNo, max(salary), min(salary)
from Employee
group by deptNo
```

Even if we know that $\max(\text{salary}) > 100.000$, the above query block is not equivalent to

```
select deptNo, max(salary), min(salary)
from Employee
where salary > 100.000
group by deptNo
```

Neither is

```
select deptNo, max(salary)
from Employee
group by deptNo
having avg(salary) > 50.000
```

equivalent to

```

select deptNo, max(salary)
from Employee
where salary < 100.000
group by deptNo
having avg(salary) < 50.000

```

even if we know that `max(salary) > 100.000`.

11.2.5 ToDo

[497]

11.3 Predicate Push-Down and Pull-Up

11.4 Eliminating Redundant Joins

11.5 Distinct Pull-Up and Push-Down

11.6 Set-Valued Attributes

In this section, we investigate the effect of query rewriting on joins involving set-valued attributes in object-relational database management systems. We show that by unnesting set-valued attributes (that are stored in an internal nested representation) prior to the actual set containment or intersection join we can improve the performance of query evaluation by an order of magnitude. By giving example query evaluation plans we show the increased possibilities for the query optimizer. This section is based on [363].

11.6.1 Introduction

The growing importance of object-relational database systems (ORDBMS) [730] has kindled a renewed interest in the efficient processing of set-valued attributes. One particular problem in this area is the joining of two relations on set-valued attributes [269, 360, 608]. Recent studies have shown that finding optimal join algorithms with set-containment predicates is very hard [98]. Nevertheless, a certain level of efficiency for joins on set-valued attributes is indispensable in practice.

Obviously, brute force evaluation via a nested-loop join is not going to be very efficient. An alternative is the introduction of special operators on the physical level of a DBMS [360, 608]. Integration of new algorithms and data structures on the physical level is problematic, however. On one hand this approach will surely result in tremendous speed-ups, but on the other hand this efficiency is purchased dearly. It is very costly to implement and integrate new algorithms robustly and reliably.

We consider an alternative approach to support set-containment and non-empty intersection join queries by compiling these join predicates away. The main idea is to unnest the set-valued attributes prior to the join. Thereby, we assume a nested internal representation [607]. This is also the underlying representation for the specific join algorithms proposed so far [360, 608]. Whereas [608] concentrates on set-containment joins, we also consider joins based on non-empty intersections. Ramasamy et al. also present a query rewrite for containment queries in [608], but on an unnested external representation, which (as shown there) exhibits very poor performance. Further, the special case of empty sets was not dealt with.

The goal of our paper is to show that by rewriting queries we can compile away the original set-containment or intersection join. As our experiments with DB2 show, our rewrite results in speed-up factors that grow linearly in the size of the input relations as compared to quadratic growth for brute-force nested-loop evaluation. The advantage of this approach—as compared to [360, 608]—is that no new join algorithms have to be added to the database system.

11.6.2 Preliminaries

In this section we give an overview of the definition of the set type. Due to the deferral of set types to SQL-4 [247], we use a syntax similar to that of Informix¹. A possible example declaration of a table with a set-valued attribute is:

```
create table ngrams (
  setID   integer not null primary key,
  content set<char(3)>
);
```

`setID` is the key of the relation, whereas `content` stores the actual set. The components of a set can be any built-in or user-defined type. In our case we used `set<char(3)>`, because we wanted to store 3-grams (see also Section ??). We further assume that on set-valued attributes the standard set operations and comparison operators are available.

Our rewriting method is based on unnesting the internal nested representation. The following view defining the unnested version of the above table keeps our representation more concise:

```
create view view_ngrams(setID, d, card) as (
  (select ngrams.setID, d.value, count(ngrams.content)
   from ngrams, table(unnest<char(3)>(ngrams.content)) d)
 union all
  (select ngrams.setID, NULL, 0)
   from ngrams
   where count(ngrams.content) = 0)
);
```

¹<http://www.informix.com/documentation/>

where `setID` identifies the corresponding set, `d` takes on the different values in `content` and `card` is the cardinality of the set. We also need `unnest<char(3)>`, a table function that returns a set in the form of a relation. As `unnest<char(3)>` returns an empty relation for an empty set, we have to consider this special case in the second subquery of the union statement, inserting a tuple containing a dummy value.

11.6.3 Query Rewrite

We are now ready to describe the queries we used to compare the nested and unnested approach. We concentrate on joins based on subset-equal and non-empty intersection predicates, because these are the difficult cases as shown in [98]. We have skipped joins involving predicates based on equality, because the efficient evaluation of these predicates is much simpler and can be done in a straightforward fashion (see [360]).

Checking Subset Equal Relation

Here is a query template for a join based on a subset-equal predicate:

```
select n_1.setID, n_2.setID
from   ngrams n_1, ngrams n_2
where  is_subseteq(n_1.content, n_2.content) <> 0;
```

(The comparison with 0 is only needed for DB2, which does not understand the type `bool`.)

This query can be rewritten as follows. The basic idea is to join the unnested version of the table based on the set elements, group the tuples by their set identifiers, count the number of elements for every set identifier and compare this number with the original counts. The filter predicate `vn1.card <= vn2.card` discards some sets that cannot be in the result of the set-containment join. We also consider the case of empty sets in the second part of the query. Summarizing the rewritten query we get

```
(select vn1.setID, vn2.setID
 from   view_ngrams vn1, view_ngrams vn2
 where  vn1.d = vn2.d
 and    vn1.card <= vn2.card
 group by vn1.setID, vn1.card, vn2.setID, vn2.card
 having count(*) = vn1.card)
union all
(select vn1.setID, vn2.setID
 from   view_ngrams vn1, view_ngrams vn2
 where  vn1.card = 0);
```

Checking Non-empty Intersection

Our query template for joins based on non-empty intersections looks as follows.


```
select n_1.setID, n_2.setID
from   ngrams n_1, ngrams n_2
where  intersects(n_1.content, n_2.content) <> 0;
```

The formulation of the unnested query is much simpler than the unnested query in Section 11.6.3. Due to our view definition, not much rewriting is necessary. We just have to take care of empty sets again, although this time in a different, simpler way.

```
select distinct vn1.setID, vn2.setID
from   view_ngrams vn1, view_ngrams vn2
where  vn1.d = vn2.d
and    vn1.card > 0;
```

11.7 Bibliography

This section is based on the investigations by Helmer and Moerkotte [363]. There, we also find a performance evaluation indicating that the rewrites depending on the relation sizes result in speed-up factors between 5 and 50 even for moderately sized relations. Nevertheless, it is argued there, that support for set-valued attributes must be build into the DBMS. A viable alternative to the rewrites presented here is the usage of special join algorithms for join predicates involving set-valued attributes [269, 359, 360, 499, 514, 515, 608]. Nevertheless, as has been shown by Cai, Chakaravarthy, Kaushik, and Naughton, dealing with set-valued attributes in joins theoretically (and of course practical) difficult issue [98]. Last, to efficiently support simple selection predicates on set-valued attributes, special index structures should be incorporated into the DBMS [361, 362, 364].

```

IU::addEqualityClassUnderThis(IU* lIU){
    IU*lRepresentativeThis = this -> getEqualityRepresentativeIU;
    IU*lRepresentativeArg = aIU -> getEqualityRepresentativeIU;

    lRepresentativeArg -> _equalityRepresentative =
    lRepresentativeThis;
    if(lRepresentativeArg -> _equalityClassRank >=
        lRepresentativeThis -> _equalityClassRank){
        lRepresentativeThis -> _equalityClassRank =
        lRepresentativeArg -> _equalityClassRank + 1;
    }
}

IU::addEqualityPredicate(Compositing* p){
    IU*lLeft = p -> leftIU;
    IU*lRight = p -> rightIU;
    if (p -> isEqualityPredicateIU &&
        lLeft -> getEqualityRepresentativeIU ==
        lRight -> getEqualityRepresentativeIU){
        if(lLeft -> isBoundToConstantIU) {
            lLeft -> addEqualityClassUnderThis(lRight);
        }else
        if(lRight -> isBoundToConstantIU){
            lRight -> addEqualityClassUnderThis(lLeft),
        }else
        if (lLeft -> _equalityClassRank > lRight ->
            _equalityClassRank){
            lLeft -> addEqualityClassUnderThis(lRight)
        }else{
            lright -> addEqualityClassUnderThis(lLeft)
        }
    }
}

IU* IU:: getEqualityRepresentativeIU(){
    if (this == _equalityRepresentative){
        _equalityRepresentative = _equalityRepresentative ->
        getEqualityRepresentativeIU;
    }
    return _equalityRepresentative;
}

```

Figure 11.3:

A1		$X \leq X$
A2	$X < Y \Rightarrow$	$X \leq Y$
A3	$X < Y \Rightarrow$	$X \neq Y$
A4	$X \leq Y \wedge X \neq Y \Rightarrow$	$X < Y$
A5	$X \neq Y \Rightarrow$	$Y \neq X$
A6	$X < Y \wedge Y < Z \Rightarrow$	$X < Z$
A7	$X \leq Y \wedge Y \leq Z \Rightarrow$	$X \leq Z$
A8	$X \leq Z \wedge Z \leq Y \wedge X \leq W \wedge W \leq Y \wedge W \neq Z \Rightarrow$	$X \neq Y$

Figure 11.4: Axioms for inequality

Chapter 12

View Merging

12.1 View Resolution

View merging can be as simple as replacing the view name in the **from** clause of a query by the view definition. We would like to call this step *view resolution*. This then results in a query with nesting in the **from** clause that can subsequently be unnested (see ??). Consider the following example: XXX Example Other examples are given below. One must be careful not to produce variable clashes. Especially if a view is referenced several times, variables must be renamed.

12.2 Simple View Merging

Of course, these two steps can be merged into one step. The overall effect is then that the view name is replaced by all the entries in the **from** clause of the view definition and the predicate contained in the **where** clause of the view definition is conjunctively added to **where** clause of the query block whose **from** clause contained the view name. Consider the following view definition

```
create view
```

which is referenced in the following query:

View merging results in

However, there are a few pitfalls. This simple version of view merging can only be applied to simple select-project-join queries not containing duplicate elimination, set operations, grouping or aggregation. In these cases, complex view merging must be applied.

12.3 Predicate Move Around (Predicate pull-up and push-down)

If unnesting is not implemented or not possible, several techniques like predicate move around, semi-join techniques and magic rewriting allow the copying of predicates from one block into another block in order to reduce the number of qualifying tuples [474, 536, 537, 538, 680].

Let us briefly illustrate the main idea by means of a simple example query

```

select  e.name
from    Employee e,
         (select  d.name, d.dno
          from    Department d
          where   d.dno = e.dno and
                d.boss.name = e.name and
                d.boss.name like '%S') as D(dname,ddno)
where   e.dno between 1 and 10

```

which can be rewritten by predicate move around to

```

select  e.name
from    Employee e,
         (select  d.name, d.dno
          from    Department d
          where   d.dno = e.dno and
                d.boss.name = e.name and
                d.dno between 1 and 10 and
                d.boss.name like '%S') as D(dname,dd no)
where   e.dno between 1 and 10 and
         e.name like '%S'

```

Predicate push-down and pull-up often occurs in conjunction with views. Let us therefore consider some examples. The following view that cannot be simply merged because it contains a **union** operator. Consider the case where there are two different employee tables that are unioned in a view.

```

create view Emp(eno, name, salary, dno) as
select  e1.eno, e1.name, e1.salary, e1.dno
from    Emp1[e1]
union all
select  e2.eno, e2.name, e2.salary, e2.dno
from    Emp2[e2]

```

Simple view merging cannot be applied to the query

```

select  e.eno, e.name
from    Emp[e]
where   e.salary > 150000

```

but view resolution with a subsequent push-down of the predicate `e.salary > 150.000` will result in

```

select e.eno, e.name
from ( select e1.eno, e1.name, e1.salary, e1.dno
        from Emp1[e1]
        where e1.salary > 150000)
union all select e2.eno, e2.name, e2.salary, e2.dno
        from Emp2[e2]
        where e2.salary > 150000)

```

Note that we did not eliminate unneeded columns/attributes. Further note that we can now exploit possible indexes on `Emp1.salary` and `Emp2.salary`. In case **union** would have been used in the view definition, the rewritten query would also contain **union** requiring a duplicate elimination.

Here is another example where pushing a predicate down results in much more efficient plans. Given the view

```

define view EmpStat as
select e.dno, min(e.salary) minSal, max(e.salary) maxSal, avg(e.salary) avgSal
from Emp[e]
group by e.dno

```

the query

```

select *
from EmpStat[e]
where e.dno = 10

```

can be rewritten to

```

select e.dno, min(e.salary) minSal, max(e.salary) maxSal, avg(e.salary) avgSal
from Emp[e]
where e.dno = 10
group by e.dno

```

which can be further simplified to

```

select e.dno, min(e.salary) minSal, max(e.salary) maxSal, avg(e.salary) avgSal
from Emp[e]
where e.dno = 10

```

12.4 Complex View Merging

12.4.1 Views with Distinct

XXX TODO views with distinct

12.4.2 Views with Group-By and Aggregation

Consider the following view with a group-by clause and aggregation:

```
create view AvgSalary as
select e.dno, avg(e.salary) as avgSalary
from Emp[e]
group by e.dno
```

The following query uses this view:

```
select d.name, s.avgSalary)
from Dept[d], AvgSalary[s]
where d.location = 'Paris' and
      d.dno = s.dno
```

Using the view definition, this query can be rewritten to

```
select d.name, avg(e.salary) as avgSalary
from Dept[d], Emp[e]
where d.location = 'Paris' and
      d.dno = e.dno
group by d.ROWID, d.name
```

where `d.ROWID` is either a key-attribute like `d.dno` or a unique row identifier of the tuples in `Dept`. Or course, this transformation is not valid in general. The primary condition here is that we have a key-foreign key join. More specifically, `d.dno` must be the key of the `Dept` table or it must be a unique attribute.

Applying simple view resolution results in:

```
select d.name, s.avgSalary)
from Dept[d], (select e.dno, avg(salary) as avgSalary
               from Emp[e]
               group by e.dno) [s]
where d.location = 'Paris' and
      d.dno = s.dno
```

This query can then be unnested using the techniques of Section ??.

Sometimes strange results occur. Consider for example the view

```
define view EmpStat as
select e.dno, min(e.salary) minSal, max(e.salary) maxSal, avg(e.salary) avgSal
from Emp[e]
group by e.dno
```

If the user issues the query

```
select avg(minSal), avg(maxSal), avg(avgSal)
from EmpStat
```

view merging results in

```
select avg(min(e.salary)), avg(max(e.salary)), avg(avg(e.salary))
from Emp[e]
group by e.dno
```

This is perfectly o.k. You just need to think twice about it. The resulting plan will contain two group operations: XXX Plan

12.4.3 Views in IN predicates

Consider a view that contains the minimum salary for each department

```
create view MinSalary as
select e.dno, min(e.salary) as minSalary
from Emp[e]
group by e.dno
```

and a query asking for all those employees together with their salaries in Parisian departments earning the minimum salary:

```
select e.name, e.salary
from Emp[e], Dept[d]
where e.dno = d.dno and
      d.location = 'Paris' and
      (e.dno, e.sal) in MinSalary
```

This query can be rewritten to:

```
select e.name, e.salary
from Emp[e], Dept[d], Emp[e2]
where e.dno = d.dno and
      d.location = 'Paris' and
      e.dno = e2.dno
group by e.ROWID, d.ROWID, e.name, e.salary
having e.salary = min(e2.sal)
```

Note that the employee relation occurs twice. Avoiding to scan the employee relation twice can be done as follows:

12.4.4 Final Remarks

Not all views can be merged. If for example a `rownum` function that numbers rows in a table is used in a view definition for a result column, then the view cannot be merged. Unmerged views will remain as nested subqueries with

two alternative evaluation strategies: Either they will be evaluated as nested queries, that is for every row produced by some outer producer the view is evaluated, or the view will be materialized into a temporary table. Whatever is more efficient must be chosen by the plan generator. However, techniques for deriving additional predicates and subsequent techniques such as predicate move around (predicate pull-down, push-down) are still applicable.

12.5 Bibliography

Chapter 13

Quantifier treatment

13.1 Pseudo-Quantifiers

Again, the clue to rewrite subqueries with a **ANY** or **ALL** predicate is to apply aggregate functions [268]. A predicate of the form

```
< ANY (select ...  
        from ...  
        where ...)
```

can be transformed into the equivalent predicate

```
< (select max(...)  
   from ...  
   where ...)
```

Analogously, a predicate of the form

```
< ALL (select ...  
        from ...  
        where ...)
```

can be transformed into the equivalent predicate

```
< (select min(...)  
   from ...  
   where ...)
```

In the above rewrite rules, the predicate $<$ can be replaced by $=$, \leq , etc. If the predicate is $>$ or \geq then the above rules are flipped. For example, a predicate of the form $>$ **ANY** becomes $>$ **select min** and $>$ **ALL** becomes $>$ **select max**.

After the rewrites have been applied, the Type A or Type JA unnesting techniques can be applied, depending on the details of the inner query block.

13.2 Existential quantifier

Existential quantifiers can be seen as special aggregate functions and query blocks exhibiting an existential quantifier can be unnested accordingly [188]. For example, an independent existential subquery can be treated the same way as a Type A query. Nested existential quantifiers with a correlation predicate can be unnested using a semi-join. Other approaches rewrite (existential) quantifiers using the aggregate function *count* [268]. Consider the partial query pattern

```

...
where exists (select ...
from          ...
where        ...)
```

It is equivalent to

```

...
where 0 > (select count(...)
from   ...
where ...)
```

A **not exists** like in

```

...
where not exists (select ...
from           ...
where         ...)
```

is equivalent to

```

...
where 0 = (select count(...)
from   ...
where ...)
```

After these rewrites have been applied, the Type A or Type JA unnesting techniques can be applied, depending on the details of the inner query block.

13.3 Universal quantifier

Universal quantification is a little more complex. An overview is provided in [156]. Here is the prototypical OQL query pattern upon which our discussion

Case-No.	1	2	3	4	5	6	7	8
	$p()$	$p()$	$p()$	$p()$	$p(e_1)$	$p(e_1)$	$p(e_1)$	$p(e_1)$
	$q()$	$q(e_1)$	$q(e_2)$	$q(e_1, e_2)$	$q()$	$q(e_1)$	$q(e_2)$	$q(e_1, e_2)$
Case-No.	9	10	11	12	13	14	15	16
	$p(e_2)$	$p(e_2)$	$p(e_2)$	$p(e_2)$	$p(e_1, e_2)$	$p(e_1, e_2)$	$p(e_1, e_2)$	$p(e_1, e_2)$
	$q()$	$q(e_1)$	$q(e_2)$	$q(e_1, e_2)$	$q()$	$q(e_1)$	$q(e_2)$	$q(e_1, e_2)$

Table 13.1: Classification Scheme According to the Variable Bindings

of universal quantifiers nested within a query block is based:

$$Q \equiv \begin{array}{l} \text{select } e_1 \\ \text{from } e_1 \text{ in } E_1 \\ \text{where for all } e_2 \text{ in select } e_2 \\ \qquad \qquad \qquad \text{from } e_2 \text{ in } E_2 \\ \qquad \qquad \qquad \text{where } p: \\ \qquad \qquad \qquad \qquad \qquad \qquad q \end{array}$$

where p (called the *range predicate*) and q (called the *quantifier predicate*) are predicates in a subset of the variables $\{e_1, e_2\}$. This query pattern is denoted by Q .

In order to emphasize the (non-)occurrence of variables in a predicate p , we write $p(e_1, \dots, e_n)$ if p depends on the variables e_1, \dots, e_n . Using this convention, we can list all the possible cases of variable occurrence. Since both e_1 and e_2 may or may not occur in p or q , we have to consider 16 cases (see Table 13.1). All cases but 12, 15, and 16 are rather trivial. Class 12 queries can be unnested by replacing the universal quantifier by a division, set difference, anti-semijoin, or counting. Class 15 queries are treated by set difference, anti-semijoin or grouping with count aggregation. For Class 16 queries, the alternatives are set difference, anti-semijoin, and grouping with count aggregation. In all cases, special care has to be taken regarding NULL values. For details see [156].

Class 12 Let us first consider an example of a Class 12 query.

```

select  al.name
from    al in Airline
where  for all ap in (select  ap
                      from    ap in Airport
                      where   apctry = 'USA'):
ap in al.lounges

```

Define $U \equiv \pi_{ap}(\sigma_{apctry='USA'}(\text{Airport}[ap, apctry]))$. Then the three alternative algebraic expressions equivalent to this query are

- plan with division:
if $U = \emptyset$

then $\text{Airline}[\text{name}]$
else $\mu_{\text{ap:lounges}}(\text{Airline}[\text{name}, \text{lounges}]) \div U$

- plan with set difference:

$$\text{Airline}[\text{name}] \setminus (\pi_{\text{name}}(U \bowtie_{\text{ap} \notin \text{lounges}} \text{Airline}[\text{name}, \text{lounges}]))$$

- plan with anti-semijoin:

$$\pi_{\text{name}}(U \overline{\bowtie}_{\text{ap} \notin \text{lounges}} \text{Airline}[\text{name}, \text{lounges}])$$

This plan is only valid, if the projected attributes of *Airline* form a superkey.

The plan with the anti-semijoin is typically the most efficient.

In general, the plan with division is [542, 301]:

$$\text{if}_{\sigma_{p(e_2)}(E_2[e_2]) \neq \emptyset}(((E_1[e_1] \bowtie_{q(e_1, e_2)} E_2[e_2]) \div \sigma_{p(e_2)}(E_2[e_2])), E_1[e_1])$$

In case the selection $\sigma_{p(e_2)}(E_2[e_2])$ yields at least a one tuple or object, we can apply the predicate p to the dividend, as in

$$\text{if}_{\sigma_{p(e_2)}(E_2[e_2]) \neq \emptyset}(((E_1[e_1] \bowtie_{q(e_1, e_2)} \sigma_{p(e_2)}(E_2[e_2])) \div \sigma_{p(e_2)}(E_2[e_2])), E_1[e_1]).$$

If the quantifier predicate $q(e_1, e_2)$ is of the form $e_2 \in e_1.\text{SetAttribute}$, then the join can be replaced by an unnest operator:

$$\text{if}_{\sigma_{p(e_2)}(E_2[e_2]) \neq \emptyset}((\mu_{e_2:\text{SetAttribute}}(E_1[e_1, \text{SetAttribute}]) \div \sigma_{p(e_2)}(E_2[e_2])), E_1[e_1])$$

Using set difference, the translation is

$$E_1[e_1] \setminus \pi_{e_1}((E_1[e_1] \times \sigma_{p(e_2)}(E_2[e_2])) \setminus (E_1[e_1] \bowtie_{q(e_1, e_2)} \sigma_{p(e_2)}(E_2[e_2])))$$

which can be optimized to

$$E_1[e_1] \setminus E_1[e_1] \bowtie_{\neg q(e_1, e_2)} \sigma_{p(e_2)}(E_2[e_2])$$

This plan is mentioned in [719], however using a regular join instead of a semi-join.

The anti-semijoin can be employed to eliminate the set difference yielding the following plan:

$$E_1[e_1] \overline{\bowtie}_{\neg q(e_1, e_2)} \sigma_{p(e_2)}(E_2[e_2])$$

This plan is in many cases the most efficient plan. However, the correctness of this plan depends on the uniqueness of e_1 , i.e., the attribute(s) e_1 must be a (super) key of E_1 . This is especially fulfilled in the object-oriented context if e_1 consists of or contains the object identifier.

We do not present the plans based group and count operations (see [156]).

Class 15 Here is an example query of Class 15:

```

select al.name
from al in Airline
where for all f in (
  select f
  from f in Flight
  where al = f.carrier):
  f.to.apctry != "Libya"

```

The quantifier's range formula $\sigma_{p(e_1, e_2)}(E_2[e_2])$ is obviously not closed. It contains the free variable e_1 . According to the reduction algorithm of Codd [169], the division plan is

$$(E_1[e_1] \bowtie_{\neg p(e_1, e_2) \vee q(e_2)} E_2[e_2]) \div E_2[e_2].$$

The plan with set difference is

$$E_1[e_1] \setminus \pi_{e_1}((E_1[e_1] \bowtie_{p(e_1, e_2)} E_2[e_2]) \setminus (E_1[e_1] \bowtie_{p(e_1, e_2)} \sigma_{q(e_2)}(E_2[e_2])))$$

and the most efficient plan using the anti-semijoin is

$$E_1[e_1] \overline{\bowtie}_{p(e_1, e_2)} \sigma_{\neg q(e_2)}(E_2[e_2]).$$

Class 16 Here is an example Class 16 query:

```

select al.name
from al in Airline
where for all ap in (
  select ap
  from ap in Airport
  where apctry = alctry):
  ap in al.lounges

```

The range predicate again depends on the outer level variable e_1 . A valid division plan looks similar to the one for Class 15. A plan with set difference is

$$E_1[e_1] \setminus \pi_{e_1}((E_1[e_1] \bowtie_{p(e_1, e_2)} E_2[e_2]) \setminus (E_1[e_1] \bowtie_{p(e_1, e_2) \wedge q(e_1, e_2)} E_2[e_2])).$$

This plan can first be refined by replacing the set difference of the two join expressions by a semijoin result in

$$E_1[e_1] \setminus (E_1[e_1] \overline{\bowtie}_{p(e_1, e_2) \wedge \neg q(e_1, e_2)} E_2[e_2])$$

Finally, the remaining set difference is transformed into an anti-semijoin which also covers the semijoin:

$$E_1[e_1] \overline{\bowtie}_{p(e_1, e_2) \wedge \neg q(e_1, e_2)} E_2[e_2].$$

Again, the uniqueness constraint on $E_2[e_2]$ is required for this most efficient plan to be valid.

For all discussed classes, problems with NULL values might occur. In that case, the plans have to be refined [156].

13.4 Bibliography

[397] [188] [156] [610, 603]

Chapter 14

Unnesting Nested Queries

The first step in unnesting a query is *view merging*. This is simply the replacement of a view name by the view definition. The result will always be a nested query. Unnesting a nested query that resulted from view merging is not different from unnesting any other nested query. However, due to a lack of orthogonality, the kinds of nesting arising from view merging can be different from that of “regular” nested queries. Several problems add to the complexity of query unnesting.

- Special cases like empty results lead easily to bugs like the famous count bug [424, 431, 268, 539, 540].
- If the nested query contains a grouping, special rules are needed to pull up grouping operators [133].
- Special care has to be taken for a correct duplicate treatment [468, 586, 681, 682].

The main reason for the problems was that SQL lacked expressiveness and unnesting took place at the query language level. The most important construct needed for correctly unnesting queries are outer-joins [188, 268, 421, 227, 539]. After their introduction into SQL and their usage for unnesting, reordering of outer-joins became an important topic [75, 188, 263, 539, 626]. Lately, a unifying framework for different unnesting strategies was proposed in [540].

14.1 Classification of nested queries

We start by extending the classification of nested queries given by Kim [424]. We restrict ourselves to a single nested block. Kim’s classification introduces five types of nested queries one of which is not used here (Type D). The four remaining types are

Type A nested queries have a constant inner block returning single elements.

Type N nested queries have a constant inner block returning sets.

Type J nested queries have an inner block that is dependent on the outer block and return a set.

Type JA nested queries have an inner block that is dependent on the outer block and return a single element.

Obviously, the need for extending the relational classification arises from the richness of the oo model compared to the relational one and its impact on the query language. The classification we propose has three dimensions: the original one plus two that are required by the following oo characteristics. In the oo context, as opposed to the relational, (i) nested blocks may be located in any clause of a **select-from-where** query and (ii) a dependency (i.e., reference to a variable of the outer block) may be expressed in any clause of a query's inner block. We restrict the presentation to queries of type A/N/J/JA with nesting and dependency (J/JA only) in the **where** clauses.

As in the relational context, the optimization of nested queries is done by unnesting, using different kinds of joins and group operators. There are two good reasons for unnesting nested queries. The first is that the underlying evaluation plan of a nested query relies on nested loops that, as shown in [424], can be very inefficient. On the other hand, we know of good algorithms for joins and group operations (using indexes, sorting, hashing). The second reason is that algebraic operators have nice properties that can be used for further rewriting whereas nested algebraic expressions don't have them a priori.

14.2 Queries of Type A

In this case, the nested query does not have any reference to any variable defined outside its block. This is equivalent to say that the nested query does not contain any free variables. This allows its independent evaluation. It can be moved outside the query block and the result produced by the nested query can be plugged in later. This way we avoid multiple evaluation of the nested query. The treatment is independent of the place where the nesting occurred.

Consider the following simple example:

```

select x1                                define m = max(select x2s
from x1 in Employee                      from x2 in Employee)
where x1.TotSales =                        define x2s = x2.TotSales
      max (select x2.TotSales              select x1
          from x2 in Employee)           from x1 in Employee
                                          where x1s = m
                                          define x1s = x1.TotSales

```

The original query is on the left-hand side. The rewritten query after NFST is shown on the right-hand side—a convention holding throughout the rest of this section. Note that the first **define** entry is independent of the second block. Hence, it is written before it. Within the implementation of the query compiler it is convenient to have an artificial outer SFWD-block to which the outer *define* clause then belongs.

Type A nested queries can be unnested by moving them one block up (like in the example). Sometimes, more efficient ways to to unnest these queries are

possible. In the example the extent of *Student* has to be scanned twice. This can be avoided by introducing the new algebraic operator *MAX* defined as

$$MAX_f(e) := \{x | x \in e, f(x) = \max_{y \in e}(f(y))\}$$

The *MAX* operator can be computed in a single pass over *e*.

Using *MAX* the above query can be expressed in the algebra as

$$q \equiv MAX_{s.age}(Student[s])$$

14.3 Queries of Type N

We discuss three different kinds of predicates occurring within the outer **where** clause:

1. $f(\vec{x})$ **in select** ...
2. **not** ($f(\vec{x})$ **in select** ...)
3. $f(\vec{x}) = (\subseteq, \supseteq, \dots)$ **select** ...

where \vec{x} represents variables of the outer block, f a function (or subquery) on these variables and $=, \subseteq, \supseteq, \dots$ are set comparisons. Other cases are possible but not necessarily unnestable.

1. Type N queries with an **in** operator can be transformed into a semi-join by using the following equivalence¹:

$$\begin{aligned} \sigma_{A_1 \in \chi_{A_2}(e_2)} e_1 &\equiv e_1 \bowtie_{A_1=A_2} e_2 \\ &\text{if } A_i \subseteq \mathcal{A}(e_i), \mathcal{F}(e_2) \cap \mathcal{A}(e_1) = \emptyset \end{aligned} \quad (14.1)$$

The first condition is obvious, the second merely stipulates that expression e_2 must be independent of expression e_1 . The interest for having this equivalence and the following is obvious. As stated previously, semi-joins and anti-joins can be implemented efficiently and they allow further rewriting. Note that special care has to be taken in order to treat duplicates correctly. We use the annotations introduced earlier, to indicate the correct treatment of duplicates.

2. Also inspired by the relational type N unnesting is the following equivalence which turns a type N query with a negated **in** operator into an anti-join:

$$\begin{aligned} \sigma_{A_1 \notin \chi_{A_2}(e_2)} e_1 &\equiv e_1 \bar{\bowtie}_{A_1=A_2} e_2 \\ &\text{if } A_i \subseteq \mathcal{A}(e_i), \mathcal{F}(e_2) \cap \mathcal{A}(e_1) = \emptyset \end{aligned} \quad (14.2)$$

We refer to [162, 719] for the third case.

These algebraic equivalences are not helpful in themselves. What is needed is a rewrite of the internal representation. We illustrate this for the first equivalence. Consider the example query

¹ \mathcal{A} is the set of defined IUs, \mathcal{F} is the set of free IUs

```

select  e.name
from    Employee e
where   e.dno in (select  d.dno
                   from    Department d
                   where   d.name like 'S%')

```

and its internal representation in textual form:

```

select  en
from    Employee[e] (J)
where   ed in dnos
define en = e.name
         ed = d.dno
         dnos = (select  dd
                 from    Department[d](J)
                 where   dn like 'S%'
                 define dd = d.dno
                       dn = d.dname)

```

The equivalence allows us to replace the **in** predicate by a join. On the internal representation, this is reflected by adding the **from** clause of the subquery to the outer **from** clause, adding the subquery's **define** clause entries to the outer query's **define** and adding the subquery's predicate to the outer query's predicate. The result is

```

select  en
from    Employee[e] (J)
         Department[d] (SJ)
where   ed = dd and
         dn like 'S%'
define en = e.name
         ed = e.dno
         dd = d.dno
         dn = d.dname

```

The equivalence allows us to replace the **in** predicate by a join. Note that we had to be careful about possible duplicates that would have been introduced if we added a simple join between departments and employees, since multiple department names might have started with a 'S'. Adding a **distinct** would not have helped since multiple employees might have the same name. Hence, the correct solution is to add a (SJ) annotation to the departments.

In the presence of key or unique constraints on the inner join column, no phantom duplicates can be produced. Consider the following example:

```

select  *
from    Account[a]
where   a.custno IN (select c.custno FROM Customer)

```

If `c.custno` is a key or there is a unique constraint on it, then no phantom duplicates can be introduced. Hence, it is safe to rewrite the query as

```
select a.*
from Account[a], Customer[c]
where a.custno = c.custno
```

If there further exists a reference constraint, we can eliminate the join altogether. For example, if the create table statement for `Account` contains a **references** `Customer` **on** `custno` constraint on `custno`, then the query is equivalent to

```
select *
from Account[a]
```

Another possibility would be to eliminate duplicates on the subqueries result. As an example consider the following query:

```
select *
from Customer[c]
where c.location in (select d.location
                    from Dept[d])
```

This query can be rewritten to

```
select *
from Customer[c],
     (select distinct d.location
      from Dept[d])
where c.location = d.location
```

Let us briefly consider queries nested in a set comparison predicate. This case does not have a counterpart in SQL. However, if we formulate the corresponding queries on a relational schema using the non-standard SQL found in [424], they would be of Type D—resolved by a division. Using standard SQL, they would require a double nesting using **EXISTS** operations. Treating Type D queries by a relational division can only treat very specific queries where the comparison predicate corresponds, in our context, to a non-strict inclusion as in the example below. The query returns the employees who have sold all the expensive products.

```
select x
from x in Employee
where x.SoldItems  $\supseteq$ 
      select i
      from i in Item
      where i.price > 20000
define ExpItems = select i
                  from i in Item
                  where p > 20000
                  define p = i.price
select x
from x in Employee
where xsi  $\supseteq$  ExpItems
define xsi = x.SoldItems
```

One solution to evaluate this query is to use a technique similar to that of [424] and add to our algebra an object division. If the set of expensive items is important, a well implemented division operation could do much compared to a nested loop evaluation. However, we voted against this operation for three reasons. The first reason is, as we stated before, that the division is based on a non-strict inclusion of the divider set. There are no more reasons to have this inclusion than any other set comparison (\subseteq , \supset , \dots). Accordingly, to be coherent, we would have to introduce one operation per set comparator (as a matter of fact, this also holds for the relational context). The second reason is that division does not have particularly nice algebraic properties that we would like to exploit. The third reason is that, since object models feature set attributes, it seems more natural to add good algorithms for dealing with selections involving set comparisons than to add new algebraic operators. Further, there already exist proposals to treat restriction predicates involving set comparison operators [721]. Thus, we prefer not to rewrite the following algebraic expression which corresponds to the translation of the above query.

$$\begin{aligned} q &\equiv \chi_x(\sigma_{xsi \supseteq ExpItems}(\chi_{xsi:x.SoldItems}(Employee[x]))) \\ ExpItems &\equiv \chi_i(\sigma_{p > 20000}(\chi_{p:i.price}(Item[i]))) \end{aligned}$$

The set *ExpItems* will be evaluated first, independently of query *q*. The result of its evaluation will be used by the selection $\sigma_{xsi \supseteq ExpItems}$ in the outer block. The selection itself can be evaluated using an algorithm similar to that of a relational division.

Note that there is no need to consider the negation of set comparisons, since it is possible to define for each set comparison an equivalent negated counterpart. Consider for example $\neg(e_1 \subseteq e_2)$ and the set comparison operator $\not\subseteq$ defined as $(e_1 \not\subseteq e_2) := (e_1 \setminus e_2 \neq \emptyset)$.

14.4 Queries of Type J

For Type J queries, we distinguish the same three cases as for Type N queries. Again, queries with **in (not in)** as the connection predicate are transformed. At the algebraic level the unnesting reads:

1.

$$\begin{aligned} \sigma_{A_1 \in \chi_{A_2}(\sigma_p(e_2))}e_1 &\equiv e_1 \triangleright_{A_1=A_2 \wedge p} e_2 & (14.3) \\ &\text{if } A_i \subseteq \mathcal{A}(e_i), \mathcal{F}(p) \subseteq \mathcal{A}(e_1 \cup e_2), \mathcal{F}(e_2) \cap \mathcal{A}(e_1) = \emptyset \end{aligned}$$

This equivalence is similar to the one used for type N queries. It just takes into account a predicate *p* relying on both *e*₁ and *e*₂ (second condition).

2.

$$\begin{aligned} \sigma_{A_1 \notin \chi_{A_2}(\sigma_p(e_2))}e_1 &\equiv e_1 \triangleright_{A_1=A_2} (e_2 \triangleright_{< p} e_1) & (14.4) \\ &\text{if } A_i \subseteq \mathcal{A}(e_i), \mathcal{F}(p) \subseteq \mathcal{A}(e_1 \cup e_2), \mathcal{F}(e_2) \cap \mathcal{A}(e_1) = \emptyset \end{aligned}$$

Type J **not in** queries cannot be translated directly using an anti-join operation: a semi-join has to be performed first.

For other cases and different unnesting possibilities see [162, 719]. The alternative unnesting strategies apply outer-joins and unary and binary grouping operations.

Now, let us consider again the case featuring a set comparison. The query below returns the employees who have sold all the items with a high-tech degree larger than the sales speciality of the employee.

<pre> select x from x in Employee where x.SoldItems \supseteq select i from i in Item where i.hTD > x.speciality </pre>	<pre> select x from x in Employee where xsi \supseteq SpecialItems define xsi = x.SoldItems xs = x.speciality SpecialItems = select i from i in Item where ihTD > xs define ihTD = i.hTD </pre>
---	---

The algebraic translation of the query is splitted for reasons of clarity:

$$\begin{aligned}
 q &\equiv \chi_x(\sigma_{xsi \supseteq \text{SpecialItems}}(q_1)) \\
 q_1 &\equiv \chi_{\text{SpecialItems}: \chi_i(\sigma_{ihTD > xs})(q_2)} \\
 q_2 &\equiv \chi_{xsi: x.SoldItems, xs: x.speciality}(\text{Employee}[x]) \\
 q_3 &\equiv \chi_{ihTD: i.hTD}(\text{Item}[i])
 \end{aligned}$$

The problem here is that the nested query is not constant. In order to unnest the query and avoid several costly scans over the set of items, we have to associate with each employee its corresponding set of special items. For this, we rely on the following equivalence:

$$\begin{aligned}
 \chi_{g: f(\sigma_{A_1 \theta A_2})(e_2)}(e_1) &\equiv e_1 \Gamma_{g; A_1 \theta A_2; f} e_2 & (14.5) \\
 &\text{if } A_i \subseteq \mathcal{A}(e_i), g \notin A_1 \cup A_2, \mathcal{F}(e_2) \cap \mathcal{A}(e_1) = \emptyset
 \end{aligned}$$

Applying this equivalence on q_1 results in

$$q \equiv \chi_x(\sigma_{xsi \supseteq \text{SpecialItems}}(q_2 \Gamma_{\text{SpecialItems}; ihTD > xs; \chi_i} q_3))$$

The binary grouping operation can be implemented by adapting standard grouping algorithms. There is another alternative to this operation that will be given in the sequel.

Two remarks. First, note that the selection with set comparator \supseteq is now evaluated between two attributes. As for type N queries, we rely on good algorithms for such selections. Second, note that the application of the equivalence did not depend on the set comparison of the predicate in the outer **where** block but on the comparison of the correlation predicate within the inner block. We will come back to this point, soon.

Eqv. 14.5 is the most general equivalence for the considered type of queries. There exist two other equivalences which deal more efficiently, using simple

grouping, with two special cases. The equivalence

$$\begin{aligned} \chi_{g:f(\sigma_{A_1=A_2}(e_2))}(e_1) &\equiv \pi_{\overline{A_2}}(e_1 \bowtie_{A_1=A_2}^{g=f(\emptyset)} (\Gamma_{g;A_2;f}(e_2))) & (14.6) \\ &\text{if } A_i \subseteq \mathcal{A}(e_i), \mathcal{F}(e_2) \cap \mathcal{A}(e_1) = \emptyset, \\ &A_1 \cap A_2 = \emptyset, g \notin \mathcal{A}(e_1) \cup \mathcal{A}(e_2) \end{aligned}$$

relies on the fact that the comparison of the correlation predicate is equality. The superscript $g = f(\emptyset)$ is the default value given when there is no element in the result of the group operation which satisfies $A_1 = A_2$ for a given element of e_1 . The equivalence

$$\begin{aligned} \chi_{g:f(\sigma_{A_1 \theta A_2}(e_2))}(e_1) &\equiv \pi_{A_1:A_2}(\Gamma_{g;A_2\theta;f}(e_2)) & (14.7) \\ &\text{if } A_i \subseteq \mathcal{A}(e_i), \mathcal{F}(e_2) \cap \mathcal{A}(e_1) = \emptyset, \\ &g \notin \mathcal{A}(e_1) \cup \mathcal{A}(e_2), \\ &e_1 = \pi_{A_1:A_2}(e_2) \text{ (this implies that } A_1 = \mathcal{A}(e_1)) \end{aligned}$$

relies on the fact that there exists a common range over the variables of the correlation predicate (third condition). We believe that these two cases are more common than the general case. We will show one application of Eqv. 14.6 in this section. In the next one, we will give an example using an equivalence derived from Eqv. 14.7.

Eqv. 14.5, 14.6, and 14.7 are not only useful for unnesting type J nested queries occurring within the **where** clause in a predicate utilizing set comparison. As already remarked above, applying these equivalence solely depends on the presence of a correlation predicate. Hence, they enable the derivation of alternative unnested expressions for the **in** and **not in** cases. To see this, consider $\sigma_{A \in e_2}(e_1) \equiv \sigma_{A \in B}(\chi_{B:e_2}(e_1))$. Further, as demonstrated in the next section, they play a major role in unnesting type JA nested queries. That is why they should be considered *the core* of unnesting nested queries in the oo context.

Further, alternative unnested evaluation plans avoiding the binary grouping operator can also be achieved by applying the following equivalence which produces an intermediate flat result and then groups it

$$\begin{aligned} \chi_{g:f(\sigma_{A'_1 \theta A'_2}(e_2))}(e_1) &\equiv \Gamma_{g;A_1;f \circ \pi_{\overline{A_1}} \circ \sigma_{A_2 \neq \perp_{A_2}}}(e_1 \bowtie_{A'_1 \theta A'_2} e_2) & (14.8) \\ &\text{if } A_i = \mathcal{A}(e_i), A'_i \subseteq A_i, g \notin A_1 \cup A_2, \mathcal{F}(e_2) \cap A_1 = \emptyset \end{aligned}$$

where \perp_A is a tuple with attributes A and null values only. Which of the Eqns. 14.5–14.8 to apply is a matter of costs.

Last, there is a variant of Eqn. 14.5 in case no selection is present:

$$\begin{aligned} \chi_{g:f(e_2)}(e_1) &\equiv e_1 \Gamma_{g;true;f} e_2 & (14.9) \\ &\text{if } g \notin \mathcal{A}(e_1) \cup \mathcal{A}(e_2), \mathcal{F}(e_2) \cap \mathcal{A}(e_1) = \emptyset \end{aligned}$$

14.5 Queries of Type JA

In the relational context, the treatment of type JA queries is radically different from that of type J, N or A. It requires joins, grouping and sometimes outer-joins

[188, 268] (remember, from the previous sections, that type N/J SQL queries required anti-joins and semi-joins). In the oo context, there is no difference between type J and type JA queries. The reason is that, in order to deal with set comparison, outer-joins and grouping operations have already been introduced to treat Type J queries and the resulting equivalences apply to type J and type JA queries [162, 719]. The grouping operators have been defined to allow the application of functions to the sets of grouped elements. This function might as well be an aggregate function. Thus, by applying Eqv. 14.5–14.8 aggregated type JA queries are treated in exactly the same manner as type J queries.

Note that, if the applied function of the unary Γ in Equivalence 14.6 is an aggregate function (as implied by type JA queries), then its right-hand side is equivalent to the generalized aggregation of [188].

14.6 Alternative locations

Nesting in the from clause Thus far we have only considered unnesting in the *where* clause. We briefly consider an example for unnesting in the *from* clause in order to illuminate the problems involved in correct duplicate handling.

Consider the following query template:

```

select  R.a, S.b
from    R, (select  s.b
              from    S s
              where  Q) as S(b)
where  P

```

This query can easily be rewritten to

```

select  R.a, S.b
from    R, S
where  P and Q

```

No problems occur with duplicates. If both queries specify **select distinct**, there is no problem either. However, in

```

select  R.a, S.b
from    R, (select  distinct s.b
              from    S s
              where  Q) as S(b)
where  P

```

duplicate removal is enforced on the inner query but not on the outer. If we just unnest in the way we did before, then both possibilities, specifying **distinct** for the outer block and not specifying it, possibly results in wrong duplicate

treatment. We need the annotations introduced earlier to avoid such complications. Another alternative is to add the keys of R to the **select** clause, specify **distinct** (i.e. remove duplicates) and then project on $R.a$ and $S.b$ without duplicate removal.

Nesting in the select clause Although nothing forbids it, type A or N nesting rarely occurs in **select** clauses. Indeed, there is not much sense in associating a constant (set or element) to each element of a set. Should that happen, we rely on the first phase of the optimization process to factor out the constant block. Thus, it will only be evaluated once.

For type J/JA queries, nesting in the **select** clause is equivalent to nesting in the **where** clause. Remember that the application of Eqns. 14.5–14.8 did not depend on the predicate in the outer **where** block but on the correlation predicate within the inner block. The same kind of correlation predicates is used when type J/JA nesting occurs in the **select** clause. We illustrate this with the following type JA query that associates to each department its number of employees.

<pre> select tuple(dept: d, emps: count(select e from e in Employee where e.dept=d)) from d in Department </pre>	<pre> select tuple(dept: d, emps: ce) from d in Department define ce = count(select e from e in Employee where ed=d define ed=e.dept) </pre>
--	---

Translating this expression into the algebra yields

$$\begin{aligned}
 q &\equiv \chi_{[dept:d, emps:ce]}(\chi_{ce:q_1}(Department[d])) \\
 q_1 &\equiv count(\chi_e(\sigma_{ed=d}(\chi_{ed:e.dept}(Employee[e]))))
 \end{aligned}$$

Eqv. 14.6 can be applied yielding:

$$\begin{aligned}
 q &\equiv \chi_{[dept:d, emps:ce]}(\pi_{ed}^{-1}(Department[d] \bowtie_{d=ed}^{ce=0} (\Gamma_{ce;ed;count \circ \chi_e}(\chi_{ed:e.dept} Employee[e])))) \\
 &\equiv \pi_{ed}^{-1}(Department[d] \bowtie_{d=ed}^{ce=0} (\Gamma_{ce;ed;count \circ \chi_e}(\chi_{ed:e.dept} Employee[e])))
 \end{aligned}$$

The zero value in the superscript $ce = 0$ corresponds to the result of the **count** function on an empty set. The transformed query can be evaluated efficiently using, for instance, a sort or an index on $Employee.dept$.

There exists one type J case where another more powerful technique can be applied: a *flatten* operation is performed on the outer block, and there is no tuple constructor within the outer block's **select** clause. As shown in [161], these queries can be optimized by pushing the *flatten* operation inside until it is applied on stored attributes; thus eliminating the nesting. For completeness, we repeat the example. The example query is

<pre> flatten(select select tuple(name:c.name,age:c.age) from c in e.children where c.age < 18) from e in employee) </pre>	<pre> flatten(select g from e in employee define ec = e.children g = select tuple(name:n,age:a) from c in ec where a < 18 define n = c.name a = c.age) </pre>
---	---

The standard translation gives

$$\begin{aligned}
 q &\equiv \text{flatten}(\chi_g(\chi_{g:e_2}(\chi_{ec:e.children}(\text{Emp}[e]))) \\
 e_2 &\equiv \chi_{[name:n,age:a]}(\sigma_{a<18}(\chi_{a:c.age,n:c.name}(ec[c])))
 \end{aligned}$$

In order to push the flatten operation inside, we have to eliminate the redundant tuple extension for the attribute g :

$$\begin{aligned}
 q &\equiv \text{flatten}(\chi_{e_2}(\chi_{ec:e.children}(\text{Emp}[e]))) \\
 e_2 &\equiv \chi_{[name:n,age:a]}(\sigma_{a<18}(\chi_{a:c.age,n:c.name}(ec[c])))
 \end{aligned}$$

Now, we know that for linear $f : \{\tau\} \rightarrow \{\tau'\}$ that

$$\text{flatten}(\chi_f(e)) = f(\text{flatten}(e)) \quad (14.10)$$

Hence,

$$\begin{aligned}
 q &\equiv \chi_{[name:n,age:a]}(\text{flatten}(\chi_{e'_2})(\chi_{ec:e.children}\text{Emp}[e])) \\
 e'_2 &\equiv \sigma_{a<18}(\chi_{a:c.age,n:c.name}(ec[c])) \\
 q &\equiv \chi_{[name:n,age:a]}(\sigma_{a<18}(\text{flatten}(\chi_{e''_2}(\chi_{ec:e.children}\text{Emp}[e])))) \\
 e''_2 &\equiv \chi_{age:c.age,n:c.name}(ec[c]) \\
 q &\equiv \chi_{[name:n,age:a]}(\sigma_{a<18}(\chi_{a:c.age,n:c.name}(\text{flatten}(\chi_{ec[c]}(\chi_{ec:e.children}\text{Emp}[e]))))) \\
 &\equiv \chi_{[name:n,age:a]}(\sigma_{a<18}(\chi_{a:c.age,n:c.name}(\text{flatten}(\chi_{e.children[c]}(\text{Emp}[e]))))) \\
 &\equiv \chi_{[name:n,age:a]}(\sigma_{a<18}(\chi_{a:c.age,n:c.name}(\text{flatten}(\chi_{children[c]}(\text{Emp})))) \\
 &\equiv \sigma_{age<18}(\chi_{[name:n,age:a]}(\chi_{[a:age,n:name]}(\text{flatten}(\chi_{children}(\text{Emp})))) \\
 &\equiv \sigma_{age<18}(\chi_{[name:name,age:age]}(\text{flatten}(\chi_{children}(\text{Emp}))))
 \end{aligned}$$

where redundant tuple constructions were eliminated in the last steps. Note that the flatten operation is now applied on stored data.

14.7 Different Kinds of Dependency

We distinguish three kinds of dependency: projection dependency (a reference to an outer variables occurs in the **select** clause), range dependency (... in the **from** clause) and predicate dependency (... in the **where** clause). Above, we studied queries with predicate dependency. In the sequel, we concentrate on optimization techniques required for range and projection dependencies.

Range dependency Consider the following query exhibiting a range dependency. It returns the set of employees having the same name than one of their children.

```

select x
from x in Employee
where x.name in select c.name
from c in x.children

```

```

select x
from x in Employee
where xn in CN
define xn = x.name
define xc = x.children
define CN = select cn
from c in xc
define cn = c.name

```

The algebraic translation is:

$$q \equiv \chi_x(\sigma_{xn \in CN}(\chi_{CN:nq}(\chi_{xn:x.name,xc:x.children}(Employee[x]))))$$

$$nq \equiv \chi_{c.name}(xc[c])$$

In terms of unnesting, there is nothing one can do. Nevertheless, the **where** clause of the above query is equivalent to the application of the path expression $x.children.name$ which passes through a set. Hence, in this case, already known optimization techniques for optimizing path expressions can be applied (see also Section 33.1).

However, there exist cases where we are able to advantageously reduce range dependencies to predicate dependencies and, hence, can unnest these queries by the above introduced techniques. The reduction relies on the existence of type extents and uses *type based rewriting* (see Section 33.1). [159, 399, 522, 524, 526]. Since it has already been described, we merely present its usage as a reduction technique useful for enabling further unnesting of range dependent subqueries. The example query is

```

select tuple (e: x.name, c: select s.customer.name
from s in = x.sales
where s.customer.city
= "Karlsruhe")
from x in Employee

```

```

select tuple (e: xn, c: SCN)
from x in Employee,
define xn = x.name
define xs = x.sales
define SCN = select scn
from s in xs
where scc = "Karlsruhe"
define sc = s.customer
define scn = sc.name
define scc = sc.city

```

Translation to the algebra yields

$$q \equiv \chi_{[e:xn,c:SCN]}(\chi_{SCN:nq}(\chi_{xn:x.name,xs:x.sales}(Employee[x])))$$

$$nq \equiv \chi_{scn}(\sigma_{scc="Karlsruhe"}(\chi_{scn:sc.name,scc:sc.city}(\chi_{sc:s.city}(xs[s]))))$$

Relying on the fact that the elements of the attribute *sales* of an employee belong to the extent of the class **Sale**, the inner block of the query can be

rewritten as

$$nq \equiv \chi_{sc.name}(\sigma_{scc="Karlsruhe"}(\chi_{sc:s.customer,sc:sc.city}(\sigma_{s \in xs}(Sale[s])))))$$

Type based rewriting can be performed again using the extent of class *Customer*. This allows us, for instance, to use indexes on *Customer.city* and *Sale.customer* to evaluate the query. However, since our goal is unnesting and not general optimization, we do not detail on this. Concerning unnesting, it is important to note that the dependency no longer specifies the range ($xs[s]$) but now represents a predicate ($\sigma_{s \in xs}$). Herewith, the algebraic expression is of the same form as one resulting from a predicate dependency. Hence, our unnesting techniques apply.

Projection dependency Queries of this kind should be rare. If they occur, they do so in two different flavors. One nice one and one nasty one. The first occurs if, within the expression forming the **select** clause, an expression occurs whose variables all depend on the outer block. Then, this expression has to be computed only once for each variable combination resulting from the evaluation of the outer block. Besides this expression, the evaluation of the inner block is independent of outer variables. Hence, it can be factored out resulting in a halfway efficient evaluation plan. The nasty case, where the expression contains variables from the outer and the inner block, requires in general the nested loop evaluation or cross product.

Remark Nothing restricts variables of the outer block to occur only at one place within the inner block. If there exist several dependencies, all the corresponding unnesting techniques can be applied alternatively. Hence, if for example a range and a predicate dependency occur, the latter should be used for unnesting if the range dependency cannot be resolved by type based rewriting.

14.8 Unnesting IN

requires that the nested query produces no duplicates. this can be enforced by introducing a duplicat elimination operation.

14.9 Further reading

For those who want to read more about unnesting in the relational context we recommend the classical papers [188, 268, 280, 421, 424, 431, 468, 539, 540, 626, 681, 682] as well as the newer ones [352, 682]. For unnesting in object-oriented database systems we recommend [161, 162, 719].

14.10 History

With his seminal paper Kim opened the area of unnesting nested queries in the relational context [424]. Very quickly it became clear that enormous perfor-

mance gains are possible by avoiding nested-loops evaluation of nested query blocks (as proposed in [36]) by unnesting them. Almost as quickly, the subtleties of unnesting became apparent. The first bugs in the original approach were detected — among them the famous count bug [431]. Retrospectively, we can detect the following problem areas:

- Special cases like empty results lead easily to bugs like the count bug [431]. These have been corrected by taking different approaches [188, 431, 421, 268, 539, 540].
- If the nested query contains a grouping, special rules are needed to pull up grouping operators [133].
- Special care has to be taken for a correct duplicate treatment [352, 468, 586, 681, 682].

The main reason for the problems was that SQL lacked expressiveness and unnesting took place at the query language level. The most important construct needed for correctly unnesting queries are outer-joins [188, 268, 421, 227, 539]. After their introduction into SQL and their usage for unnesting, reordering of outer-joins became an important topic [75, 188, 263, 539, 626]. Lately, a unifying framework for different unnesting strategies was proposed in [540].

With the advent of object-oriented databases and their query languages, unnesting once again drew some attention from the query optimization research community [161, 162, 720, 721, 719, 722]. Different from the relational unnesting strategies which mainly performed at the (extended) SQL source level, researchers preferred to use algebras that allowed nesting. For example, a predicate of a selection operator could again contain algebraic operators. Unnesting then took place at the algebraic level. The advantage of this approach are (1) it is mainly query language independent and (2) by using algebraic equivalences, correctness proofs could be delivered.

With the arrival of XQuery [240], the field was reopened by Paparizos et al. [572]. Within their approach, unnesting takes place at the algebraic level. The underlying algebra's data model is based on sets of ordered labeled trees [393]. However, instead of using a simple equivalence, a verbal description of more than one page is used to describe detection of applicability and the according unnested plan. Since this description is verbal, it is not rigorous and indeed buggy.

14.11 Bibliography

[256] [574] [612] [41] [100]

14.12 ToDo

betrachte unnesting query:

```
select [a:a,b:b]
from a in A
      b in A.b
```

ExtremumSpecialCases
Nested Relational Approach: [100]

ToDo

Chapter 15

Optimizing Queries with Materialized Views

15.1 Conjunctive Views

15.2 Views with Grouping and Aggregation

15.3 Views with Disjunction

15.4 Bibliography

materialized view with aggregates: [717],
materialized view with disjunction: [11],
SQL Server: [282]
other: [12, 125, 126, 139, 472, 732, 764, 827] [114, 118, 140, 128, 244, 411, 457, 575, 602, 673]
some more including maintenance etc: [10, 14, 47, 81, 125, 133, 175, 328, 342] [373, 410, 471, 599, 639, 228, 717] [732, 741, 740, 849, 828] [6, 215, 216, 350]
Overview: [334]
[473]
performance eval: [79]
Stacked views: [193]
recursion: [218]
with patterns (integration): [602], [217, 219], [200]

Chapter 16

Semantic Query Rewrite

16.1 Constraints and their impact on query optimization

Using Constraints: [279, 321]

16.2 Semantic Query Rewrite

Semantic query rewrite exploits knowledge (semantic information) about the content of the object base. This knowledge is typically specified by the user. We already saw one example of user-supplied information: *inverse relationships*. As we already saw, inverse relationships can be exploited for more efficient query evaluation.

Another important piece of information is knowledge about keys. In conjunction with type inference, this information can be used during query rewrite to speed up query execution. A typical example is the following query

```
select distinct *  
from           Professor p1, Professor p2  
where         p1.university.name = p2.university.name
```

By type inference, we can conclude that the expressions *p1.university* and *p2.university* are of type University. If we further knew that the name of universities are unique, that is the name is a candidate key for universities, then the query could be simplified to

```
select distinct *  
from           Professor p1, Professor p2  
where         p1.university = p2.university
```

Evaluating this query does no longer necessitate accessing the universities to retrieve their *name*.

Some systems consider even more general knowledge in form of equivalences holding over user-defined functions [1, 245]. These equivalences are then used to rewrite the query. Thereby, alternatives are generated all of which are subsequently optimized.

Semantic Query Optimization: [114]

16.3 Exploiting Uniqueness in Query Optimization

[579]

16.4 Bibliography

[73] [65] [820] Foreign functions semantic rules rewrite: [128] Conjunctive Queries, Branch Minimization: [633]

Part IV

Plan Generation

Chapter 17

Current Search Space and Its Limits

17.1 Plans with Outer Joins, Semijoins and Anti-joins

outer join reordering [255, 254, 626, 263], outer join/antijoin plan generation [611], semijoin reducer [725],

17.2 Expensive Predicates and Functions

17.3 Techniques to Reduce the Search Space

- join single row tables first
- push down SARGable predicates
- For large join queries do not apply transitivity of equality to derive new predicates and disable cross products and possibly bushy trees.

17.4 Bibliography

Chapter 18

Optimizing Queries with Disjunctions

18.1 Introduction

Simple rewrites as indicated in Section ?? for IN and OR predicates that boil down to comparisons of a column with a set of constants can eliminate disjunction from the plan or push it into a multirange index access.

Another possibility that can be used for disjunctions on single columns is to use DISJOINT UNION of plans. This is a special form of UNION where conditions ensure that no phantom duplicates are produced. The DISJOINT UNION operator merely concatenates the result tables without any further overhead like duplicate elimination.

For example a predicate of the form $x = c_1$ or $y = c_2$ where x and y are columns of the same table results in two predicates

1. $x = c_1$
2. $x \neq c_1$ AND $y = c_2$

Obviously, no row can satisfy both conditions. Hence, the query `select * from R where x = c1 or y = c2` can be safely rewritten to

```
(select * from R where x = c1) DISJOINT UNION (select * from
R where x <> c1 AND y = c2)
```

In case there are indexes on x and y efficient plans do exist. If they don't the table R needs to be scanned twice. This problem is avoided by using bypass plans.

DISJOIN UNIONS can also be used for join predicates. Consider the following example query: `select * from R, S where R.a = S.a OR R.b = S.a` This query can be rewritten to `(select * from R, S where R.a = S.a) DISJOINT UNION (select * from R, S where R.a <> S.a and R.b = S.b)` The general condition here is that all equality predicates have one side identical. Note that both tables are scanned and joined twice. Bypass plans will eliminate this problem.




Figure 18.2: CNF plans

CNF plans never produce duplicates. The evaluation of the boolean factors can stop as soon as some predicate evaluates to *true*. Again, some (expensive) predicates might be evaluated more than once in CNF plans. Figure 18.3 shows some bypass plans. Note the different output streams. It should be obvious, that a bypass plan can be more efficient than both a CNF or DNF plan. It




Figure 18.3: Bypass plans

is possible to extend the idea of bypass plans to join operators. However, this and the algorithm to generate bypass plans is beyond the scope of the current paper (see [418, 724, 158]).

18.4 Implementation remarks

The internal representation of execution plans during plan generation typically differs from that used in Rewrite I. The reason is that many plans have to be generated and space efficiency is a major concern. As in the query representation discussed earlier, the physical algebraic operators can be organized into a hierarchy. Besides their arguments, they possibly contain backpointers to the original query representation (e.g. for predicates). Sharing is a must for plan generation. Hence, subplans are heavily shared. The plan nodes are enhanced by so-called property vectors. These contain information about the plan:

- logical information
 - the set of relations joined
 - the set of predicates applied so far
 - the set of IUs computed so far
 - order information
- physical information
 - costs
 - cardinality information

For fast processing, the first three set-valued items in the logical information block are represented as bit-vectors. However, the problem is that an upper bound on the size of these bitvectors is not reasonable. Hence, they are of variant size. It is recommendable, to have a plan node factory that generates plan nodes of different length such that the bit-vectors are included in the plan node. A special interpreter class then knows the offsets and lengths of the different bitvectors and supplies the operations needed to deal with them. This bit-vector interpreter can be attached to the plan generator's control block as indicated in Fig. 25.3.

18.5 Other plan generators/query optimizer

There are plenty of other query optimizers described in the literatur. Some of my personal favorites not mentioned so far are the Blackboard query optimizer [417], the Epoq optimizer [525, 523], the Genesis optimizer [50, 55], the Gral query optimizer [60], the Lanzelotte query optimizer [451, 452, 453], the Orion optimizer [45, 46, 423], the Postgres optimizer [400, 344, 342, 343], the Prima optimizer [348, 346], the Probe optimizer [191, 190, 555], the Straube optimizer [739, 769, 736, 737, 735, 738]. Highly recommended is a description of the DB2 query optimizer(s) [275].

Also interesting to read is the first proposal for a rule-based query optimizer called Squirrel [714] and other proposals for rule-based query optimizers [251, 668, 415, 414, 498].

18.6 Bibliography

Disjunctive queries: P. Ciaccia and M. Scalas: Optimization Strategy for Relational Queries. *IEEE Transaction on Software Engineering* 15 (10), pp 1217-1235, 1989.

Kristofer Vorwerk, G. N. Paulley: On Implicate Discovery and Query Optimization. *International Database Engineering and Applications Symposium (IDEAS'02)*

Jack Minker, Rita G. Minker: Optimization of Boolean Expressions-Historical Developments. *IEEE Annals of the History of Computing* 2 (3), pp 227-238, 1980.

Chaudhuri: *SIGMOD 03*: [123]

Conjunctive Queries, Branch Minimization: [633]

Also Boolean Difference Calculus (?): [715]

Chapter 19

Grouping and Aggregation

19.1 Introduction

In general, join and grouping operations are not reorderable. Consider the following relations R and S

R	A	B	S	A	C
	a	5		a	7
	a	6		a	8

Joining these relations R and S results in

$R \bowtie S$	A	B	C
	a	5	7
	a	5	8
	a	6	7
	a	6	8

Applying $\Gamma_{A;count(*)}$ to R and $R \bowtie S$ yields

$\Gamma_{A;count(*)}(R)$	A	count (*)	$\Gamma_{A;count(*)}(R \bowtie S)$	A	count (*)
	a	2		a	4

Compare this to the result of $\Gamma_{A;count(*)}(R) \bowtie S$:

$\Gamma_{A;count(*)}(R) \bowtie S$	A	count (*)	C
	a	2	7
	a	2	8

Hence $\Gamma_{A;count(*)}(R) \bowtie S \neq \Gamma_{A;count(*)}(R \bowtie S)$.

Since grouping and join operations are in general not reorderable, it is important that a query language determines the order of grouping and join operators properly. In SQL, the grouping operator is applied after the join operators of a query block.

For example, given the relations schemata

Emp (eid, name, age, salary) and

Sold (sid, eid, date, product_id, price)

and the query

```

select    e.eid, sum (s.price) as amount
from      Emp e, Sold s
where     e.eid = s.eid and
           s.date between "2000-01-01" and "2000-12-31"
group by s.eid, s.name

```

results in the algebraic expression

$$\Pi_{e.eid, amount} (\Gamma_{s.eid; amount: \text{sum}(s.price)} (Emp[e] \bowtie_{e.eid=s.eid} \sigma_p (Sold[s])))$$

where p denotes

$$s.date \geq '2000-01-01' \wedge s.date \leq '2000-12-31'$$

Figure 20.1 (a) shows this plan graphically. Note that the grouping operator is executed last in the plan. This is the standard translation technique applied to SQL. However, Yan and Larson discovered that under certain circumstances grouping and join can be reordered [823]. In sequel work, Yan and Larson as well as Chaudhuri and Shim extended the possibilities to group before or after a join [129, 823, 824, 825, 826]. These extensions of the search space are the topic of this chapter.

Before we delve into details, let us consider an alternative plan for the above query. Here, we push down the grouping operator: The plan then becomes:

$$\Pi_{e.eid, amount} (Emp[e] \bowtie_{e.eid=s.eid} (\Gamma_{s.eid; amount: \text{sum}(s.price)} (\sigma_p (Sold[s])))$$

This plan (see also Figure 20.1 (b)) is equivalent to the former plan. Moreover, if the grouping operator strongly reduces the cardinality of

$$\sigma_{s.date \geq \dots} (Sold[s])$$

because every employee sells many items, then the latter plan might become cheaper since the join inputs are smaller than in the former plan. This motivates the search for conditions under which join and grouping operators can be reordered. Several papers discuss this reorderability and other kinds of search space extensions [129, 823, 824, 825, 826]. We will summarize their results in subsequent sections. Before that, we will take a look at aggregation functions and their properties.

The plan for the rest of the chapter is the following. First, we take a closer

todo

19.2 Aggregate Functions

SQL and many other query languages support at least five aggregation functions. These are *min*, *max*, *count*, *sum*, and *avg*. In addition, SQL allows

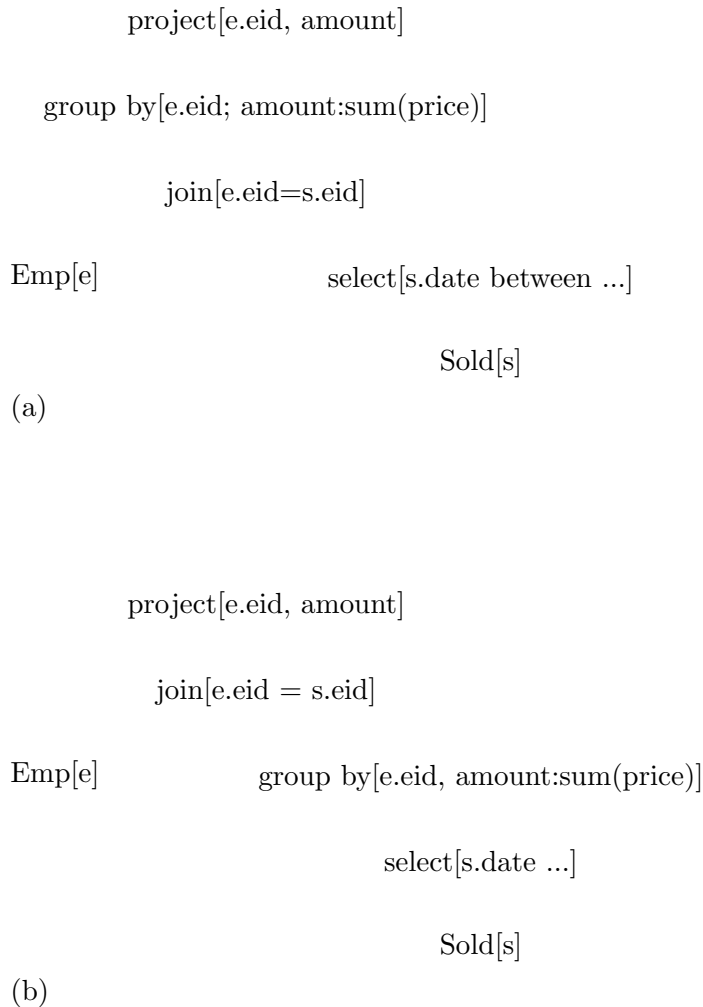


Figure 19.1: Two equivalent plans

to qualify whether duplicates are removed before computing the aggregate or whether they are also considered by the aggregation function. For example, we may specify **sum(distinct a)** or **sum(all a)** for some attribute a . The term **sum(a)** is equivalent to **sum(all a)**. From this follows that aggregation functions can be applied to sets or bags. Other query languages (OQL and XQuery) also allow lists as arguments to aggregation functions. Additionally, OQL allows arrays. Hence, aggregation functions should be defined for any bulk type.

Most query languages provide a special null value (SQL provides **NULL**, OQL **UNKNOWN**). Typically, aggregation functions can safely ignore null values. The only exception is *count*, all input is counted independent of whether its value is null or not. If we want to count only non-null values, we could imagine an expression of the form **count(distinct not null a)**. Unfortunately, this is not valid SQL. However, we will make use of a variant of *count*, that does not count null values. Let us denote this function by $count^{nn}$.

Let \mathcal{N} denote either a numeral data type (e.g. *integer* or *float*) or a tuple $[a_1 : \tau_1, \dots, a_n : \tau_n]$ where each type τ_n is a numeral data type. Further, let \mathcal{N} contain a special value *NULL* denoted by *NULL*.

A *scalar* aggregation function agg is a function with signature

$$\text{agg} : \text{bulk}(()\tau) \rightarrow \mathcal{N}$$

A scalar aggregation function $\text{agg} : \text{bulk}(()\tau) \rightarrow \mathcal{N}$ is called *decomposable* if there exist functions

$$\begin{aligned} \text{agg}^I &: \text{bulk}(()\tau) \rightarrow \mathcal{N}' \\ \text{agg}^O &: \text{bulk}(()\mathcal{N}') \rightarrow \mathcal{N} \end{aligned}$$

with

$$\text{agg}(Z) = \text{agg}^O(\text{bulk}(()\text{agg}^I(X), \text{agg}^I(Y)))$$

for all X and Y (not empty) with $Z = X \cup_{\text{Bulk}} Y$. This condition assures that $\text{agg}(Z)$ can be computed on arbitrary subsets (-lists, -bags) of Z independently and the (partial) results can be joined to yield the correct (total) result. If the condition holds, we say that agg is *decomposable* with *inner* agg^I and *outer* agg^O . Decomposability will also be applied to vectors of aggregate functions.

A decomposable scalar aggregation function $\text{agg} : \text{bulk}(()\tau) \rightarrow \mathcal{N}$ is called *reversible* if for agg^O there exists a function $(\text{agg}^O)^{-1} : \mathcal{N}', \mathcal{N}' \rightarrow \mathcal{N}'$ with

$$\text{agg}(X) = \gamma((\text{agg}^O)^{-1}(\text{agg}^I(Z), \text{agg}^I(Y)))$$

for all X , Y , and Z with $Z = X \cup_{\text{Bulk}} Y$. This condition assures that we can compute $\text{agg}(X)$ for a subset (-list, -bag) X of Z by “subtracting” its aggregated complement Y from the “total” $\text{agg}^O(\text{agg}^I(Z))$ by using $(\text{agg}^O)^{-1}$.

The fact that scalar aggregation functions can be decomposable and reversible is the basic observation upon which the efficient evaluation of aggregation functions builds.

As an example consider the scalar aggregation $\text{avg} : \{\{\{\}\}[a : \text{float}]\} \rightarrow \text{float}$ averaging the values of the attributes a of a bag of tuples with a single attribute a . It is reversible with

$$\begin{aligned} \text{agg}^I &: \{\{a : \text{float}\}\} \rightarrow [\text{sum} : \text{float}, \text{count} : \text{float}] \\ \text{agg}^O &: [\text{sum} : \text{float}, \text{count} : \text{float}], [\text{sum} : \text{float}, \text{count} : \text{float}] \rightarrow [\text{sum} : \text{float}, \text{count} : \text{float}] \\ (\text{agg}^O)^{-1} &: [\text{sum} : \text{float}, \text{count} : \text{float}], [\text{sum} : \text{float}, \text{count} : \text{float}] \rightarrow [\text{sum} : \text{float}, \text{count} : \text{float}] \\ \gamma &: [\text{sum} : \text{float}, \text{count} : \text{float}] \rightarrow \text{float} \end{aligned}$$

where

$$\begin{aligned} \text{agg}^I(X) &= [\text{sum} : \text{sum}(X.a), \text{count} : |X|] \\ \text{agg}^O([\text{sum} : s_1, \text{count} : c_1], [\text{sum} : s_2, \text{count} : c_2]) &= [\text{sum} : s_1 + s_2, \text{count} : c_1 + c_2] \\ (\text{agg}^O)^{-1}([\text{sum} : s_1, \text{count} : c_1], [\text{sum} : s_2, \text{count} : c_2]) &= [\text{sum} : s_1 - s_2, \text{count} : c_1 - c_2] \\ \gamma([\text{sum} : s, \text{count} : c]) &= s/c \end{aligned}$$

$sum(X.a)$ denotes the sum of all values of attribute a of the tuples in X , and $|X|$ denotes the cardinality of X . Note that $agg^I(\emptyset) = [sum : 0, count : 0]$, and $\gamma([sum : 0, count : 0])$ is undefined as is $avg(\emptyset)$.

aggregation
function!duplicated
duplicated
aggregation
function

Not all aggregation functions are decomposable and reversible. For instance, min and max are decomposable but not reversible. If an aggregation function is applied to a bag that has to be converted to a set, then decomposability is jeopardized for sum and $count$. That is, in SQL **sum(distinct)** and **count(distinct)** are not decomposable.

Let us look at the decomposition of our five aggregation functions. We can decompose them as follows:

$$\begin{aligned} \min(X \cup_{Bulk} Y) &= \min(\min(X), \min(Y)) \\ \max(X \cup_{Bulk} Y) &= \max(\max(X), \max(Y)) \\ \text{count}(X \cup_{Bulk} Y) &= \text{sum}(\text{count}(X), \text{count}(Y)) \\ \text{sum}(X \cup_{Bulk} Y) &= \text{sum}(\text{sum}(X), \text{sum}(Y)) \end{aligned}$$

Treatment of avg is slightly more complicated, as we have already seen above. In the presence of null values, avg is defined as $avg(X) = \text{sum}(X) / \text{count}^{NN}(X)$. Hence, we can decompose it on the basis of

$$avg(X \cup_{Bulk} Y) = \text{sum}(\text{sum}(X), \text{sum}(Y)) / (\text{count}^{NN}(X) + \text{count}^{NN}(Y))$$

It is also useful to classify aggregation functions as follows [825]:

- Class C Aggregate Functions: sum , $count$
- Class D Aggregate Functions: $sum(\text{distinct})$, $count(\text{distinct})$, min , max , $avg(\text{distinct})$

Class C aggregation functions require a multiplication by the count of the inner query block. Therefore, we sometimes need to replace the vector of aggregation functions

$$F = [a_1 : agg_1(e_1), \dots, a_n : agg_n(e_n)]$$

by the vector

$$F^{*c} = [a_1 : agg_1^{*c}(e_1), \dots, a_n : agg_n^{*c}(e_n)]$$

$agg_i^{*c}(e_i)$ is defined as $agg_i(e_i)$ if agg_i is a class D aggregate function. If agg_i is a Class C aggregate function, we define $agg_i^{*c}(e_i)$ as $agg_i(e_i) * c$ where c is a special attribute that contains the result of some count in a subquery block. F^{*c} is called *duplicated aggregation function of F*.

19.3 Normalization and Translation

19.3.1 Grouping and Aggregation in Different Query Languages

Conversion from bags to sets must be explicitly specified within the query language. As we have seen, in SQL this is done by specifying **distinct** directly after the parenthesis following the name of the aggregation function. In

OQL, the conversion function *distinct* is used. For example, the OQL query $avg(distinct(bag(1,1,2,3,3)))$ return 2. Similarly for XQuery.

We now come to an essential difference between SQL and OQL/XQuery. SQL allows expressions of the form **sum**(*a*) where *a* is a single-valued attribute. Since aggregation functions take bulk types as arguments, this expression may seem to contain a type error. Let us call this *false aggregates*. There are two cases to consider depending on whether the block where the aggregation function(s) occur exhibits a **group by** or not. . . .

19.3.2 Extensions to Algebra

19.3.3 Normalization

19.3.4 Translation

19.4 Lazy and eager group by

Lazy group by pulls a group operator up over a join operator [823, 824, 825, 826]. Eager group by does the opposite. This may also be called *Push-Down Grouping* and Pull-Up Grouping.

Consider the query:

```

select[all | distinct]    $A, \overrightarrow{F}(B)$ 
from                      $R, S$ 
where                     $p_R \wedge p_S \wedge p_{R,S}$ 
group by                  $G$ 

```

with

$$G = G_R \cup G_S, G_R \subseteq \mathcal{A}(R), G_S \subseteq \mathcal{A}(S),$$

$$\mathcal{F}(p_R) \subseteq \mathcal{A}(R), \mathcal{F}(p_S) \subseteq \mathcal{A}(S)$$

$$\mathcal{F}(p_{R,S}) \subseteq \mathcal{A}(R) \cup \mathcal{A}(S)$$

$$B \subseteq \mathcal{A}(R) \quad A = A_R \cup A_S, A_R \subseteq G_R, A_S \subseteq G_S$$

$$\text{agg}_R^I = G_R \cup \mathcal{F}(p_{R,S}) \setminus \mathcal{A}(S) \quad \kappa_R \text{ key of } R$$

$$\text{agg}_S^I = G_S \cup \mathcal{F}(p_{R,S}) \setminus \mathcal{A}(R) \quad \kappa_S \text{ key of } S$$

We are interested in the conditions under which the query can be rewritten into

```

select[all | distinct]   $A, FB$ 
from                     $R', S'$ 
where                    $p_{R,S}$ 

with

 $R'(\text{agg}_R^I, FB) \equiv$ 

select all  $\text{agg}_R^I, \overrightarrow{F}(B)$  as  $FB$ 
from       $R$ 
where     $p_R$ 
group by  $\text{agg}_R^I$ 

and

 $S'(\text{agg}_S^I) \equiv$ 

select all  $\text{agg}_R^I$ 
from       $S$ 
where     $p_S$ 

```

The following equivalence expresses this rewrite in algebraic terms.

$$\Pi_{A,F}^{[d]} \left(\Gamma_{G;F:\overrightarrow{F}(B)} \left(\sigma_{p_R}(R) \bowtie_{p_{R,S}} \sigma_{p_S}(S) \right) \right) \equiv$$

$$\Pi_{A,F}^{[d]} \left(\Gamma_{\text{agg}_R^I;F:\overrightarrow{F}(B)} \left(\sigma_{p_R}(R) \right) \bowtie_{p_{R,S}} \sigma_{p_S}(S) \right)$$

holds iff in $\sigma_{p_R \wedge p_S \wedge p_{R,S}}(R \times S)$ the following functional dependencies hold:

$FD_1 : G \rightarrow \text{agg}_R^I$

$FD_2 : \text{agg}_R^I, G_S \rightarrow \kappa_S$

Note that since $G_S \subseteq G$, this implies $G \rightarrow \kappa_S$.

FD_2 implies that for any group there is at most one join partner in S . Hence, each tuple in $\Gamma_{\text{agg}_R^I;F:\overrightarrow{F}(B)}(\sigma_{p_R}(R))$ contributes at most one row to the overall result.

FD_1 ensures that each group of the expression on the left-hand side corresponds to at most one group of the group expression on the right-hand side.

We now consider queries with a **having** clause.

In addition to the assumptions above, we have that the tables in the **from** clause can be partitioned into R and S such that R contains all aggregated columns of both the **select** and the **having** clause. We further assume that conjunctive terms in the **having** clause that do not contain aggregation functions have been moved to the **where** clause.

Let the predicate of the **having** clause have the form $H_R \wedge H_0$ where $H_R \subseteq \mathcal{A}(R)$ and $H_0 \subseteq R \cup S$ where H_0 only contains non-aggregated columns from S .

We now consider all queries of the form

select [all distinct]	$A, \vec{F}(B)$
from	R, S
where	$p_R \wedge p_S \wedge p_{R,S}$
group by	G
having	$H_0 \left(\vec{F}_0(B) \right) \wedge H_R \left(\vec{F}_R(B) \right)$

where \vec{F}_0 and \vec{F}_R are vectors of aggregation functions on the aggregated columns B .

An alternative way to express such a query is

select [all distinct]	G, FB
from	R', S
where	$c_S \wedge c_{R,S} \wedge H_0(F_0B)$
where	$R' \left(\text{agg}_R^I, FB, F_0B \right) \equiv$
select all	$\text{agg}_R^I, \vec{F}(B) \text{ as } FB, \vec{F}_0(B) \text{ as } F_0B$
from	R
where	c_R
group by	agg_R^I
having	$H_R \left(\vec{F}_R(B) \right)$

The according equivalence is [825]:

$$\begin{aligned} & \Pi_{G,F} \left(\sigma_{H_R \wedge H_0} \left(\Gamma_{G;F:\vec{F}(B),F_R:\vec{F}_R(B),F_0:\vec{F}_0(B)} \left(\sigma_{p_R \wedge p_S \wedge p_{R,S}} (R \times S) \right) \right) \right) \\ & \equiv \\ & \Pi_{G,F} \left(\sigma_{p_{R,S} \wedge p_S \wedge H_0(F_0)} \right) \left(\Pi_{G,F,F_0} \left(\sigma_{H_R} \left(\Gamma_{G;F:\vec{F}(B),F_R:\vec{F}_R(B),F_0:\vec{F}_0(B)} (R) \right) \right) \times S \right) \end{aligned}$$

19.5 Coalescing Grouping

In this section we introduce *coalescing grouping* which slightly generalizes *simple coalescing grouping* as introduced in [129].

We first illustrate the main idea by means of an example.

Given two relation schemes

Sales (pid, deptid, total_price)
 Department (did, name, region)

the query

select	region, sum (total_price) as s
from	Sales, Department
where	did = deptid
group by	region

is straightforwardly translated into the following algebraic expression:

$$\Gamma_{region;s:sum(total_price)}(\text{Sales} \bowtie_{deptid=did} \text{Department})$$

Note that Equivalence ?? cannot be applied here. However, if there are many sales performed by a department, it might be worth reducing the cardinality of the left join input by introducing an additional group operator. The result is

$$\Gamma_{region;s=sum(s')}(\Gamma_{deptid;s':sum(total_price)}(\text{Sales}) \bowtie_{deptid=did} \text{Department})$$

Note that we must keep the outer grouping.

That is, we introduced an additional group operator to reduce the cardinality of sales. This way, all subsequent joins (only one in this case) become cheaper and the additional group operator may result in a better plan.

We have the following restrictions for this section:

1. There are no NULL-values allowed for attributes occurring in the query.
2. All queries are of the form **select all**.
That is **select distinct** is not allowed.
3. All aggregation functions agg must fulfill $agg_{s_1} \cup s_2 = agg\{agg(s_1), agg(s_2)\}$ for bags s_1 and s_2 .
This has two consequences:
 - Allowed are only sum, min, max. Not allowed are avg and count.
 - For any allowed aggregation function we only allow for **agg(all ...)**.
Forbidden is **agg(distinct ...)**.
4. The query is a single-block conjunctive query with no **having** and no **order by** clause.

The above transformation is an application of the following equivalence, where R_1 and R_2 can be arbitrary algebraic expressions:

$$\Gamma_{G;A}(R_1 \bowtie_p R_2) \equiv \Gamma_{G;A_2}(\Gamma_{G_1;A_1}(R_1) \bowtie_p R_2) \quad (19.1)$$

with

$$\begin{aligned} A &= A_1 : agg_1(e_1), \dots, A_n : agg_n(e_n) \\ A_1 &= A_1^1 : agg_1^1(e_1), \dots, A_n^1 : agg_n^1(e_n) \\ A_2 &= A_1 : agg_1^2(A_1^1), \dots, A_n : agg_n^2(A_n^2) \end{aligned}$$

$$G_1 = (\mathcal{F}(p) \cup G) \cap \mathcal{A}(R_1)$$

Further, the following condition must hold for all $i(1 \leq i \leq n)$:

$$agg_i\left(\bigcup_k S_k\right) = agg_i^2\left(\bigcup_k \{agg_i^1(S_i)\}\right)$$

In the above example, we had $agg_1 = agg_1^1 = agg_1^2 = \mathbf{sum}$.

We now prove the correctness of Equivalence 20.1.

Proof:

First, note that

$$R_1 \bowtie_p R_2 = \bigcup_{t_2 \in R_2} R_1 \bowtie_p \{t_2\} \quad (19.2)$$

Second, note that for a given t_2

$$\begin{aligned} \Gamma_{G;A}(R_1[t_1]) \bowtie_p \{t_2\} &= \sigma_{p(t_1 \circ t_2)}(\Gamma_{G;A}(R_1[t_1])) \\ &= \Gamma_{G;A}(\sigma_{p(t_1 \circ t_2)}(R_1[t_1])) \\ &= \Gamma_{G;A}(R_1[t_1] \bowtie_p \{t_2\}) \end{aligned} \quad (19.3)$$

holds where we have been a little sloppy with t_1 . Applying (20.2) and (20.3) to $\Gamma_{G_1;A_1}(R_1) \bowtie_p R_2$, the inner part of the right-hand side of the equivalence yields:

$$\begin{aligned} \Gamma_{G_1;A_1}(R_1) \bowtie_p R_2 &= \bigcup_{t_2 \in R_2} \Gamma_{G_1;A_1}(R_1) \bowtie_p \{t_2\} \\ &= \bigcup_{t_2 \in R_2} \Gamma_{G_1;A_1}(R_1 \bowtie_p \{t_2\}) \end{aligned} \quad (19.4)$$

Call the last expression X.

Then the right-hand side of our equivalence becomes

$$\begin{aligned} \Gamma_{G;A_2}(X) &= \{t \circ a_2 \mid t \in \Pi_G(X), a_2 = (A_1 : a_1^2, \dots, A_n : a_n^2), \\ &\quad a_i^2 = agg_i^2(\{s.A_i^1 \mid s \in X, S|_G = t\})\} \end{aligned} \quad (19.5)$$

Applying (20.2) to the left-hand side of the equivalence yields:

$$\Gamma_{G;A}(R_1 \bowtie_p R_2) = \Gamma_{G;A} \left(\bigcup_{t_2 \in R_2} \{t_1 \circ t_2 \mid t_1 \in R_1, p(t_1 \circ t_2)\} \right) \quad (19.6)$$

Abbreviate $\bigcup_{t_2 \in R_2} \{t_1 \circ t_2 \mid t_1 \in R_1, p(t_1 \circ t_2)\}$ by Y.

Applying the definition of $\Gamma_{G;A}$ yields:

$$\begin{aligned} \{t \circ a \mid t \in \Pi_G(Y), a = (A_1 : e_1, \dots, A_n : e_n), \\ a_i = agg_i(\{e_i(s) \mid s \in Y, S|_G = t\})\} \end{aligned} \quad (19.7)$$

Compare (20.5) and (20.7). Since $\Pi_G(X) = \Pi_G(Y)$, they can only differ in their values of A_i .

Hence, it suffices to prove that $a_i^2 = a_i$ for $1 \leq i \leq n$ for any given t .

$$\begin{aligned}
a_i^2 &= \text{agg}_i^2(\{s.A_i^1 | s \in X, S|_G = t\}) \\
&= \text{agg}_i^2(\{s.A_i^1 | s \in \bigcup_{t_2 \in R_2} \Gamma_{G_1; A_1}(R_1 \bowtie_p \{t_2\}), S|_G = t\}) \\
&= \text{agg}_i^2(\{s.A_i^1 | s \in \bigcup_{t_2 \in R_2} \{t_1 \circ t_2 \circ a_1 | t_1 \in \Pi_{G_1}(R_1), p(t_1 \circ t_2), \\
&\quad a_1 = (A_1^1 : a_1^1, \dots, A_n^1 : a_n^1) \\
&\quad a_i^1 = \text{agg}_i^1(\{e_i(s_1 \circ t_2) | s_1 \in R_1, S_1|_{G_1=t_1}, p(s_1, t_2)\}) \\
&\quad S|_G = t\}) \\
&= \text{agg}_i^2(\bigcup_{t_2 \in R_2} \{\text{agg}_i^1(\{e_i(s_1 \circ t_2) | t_1 \in \Pi_{G_1}(R_1), p(t_1 \circ t_2), s_1 \in R_1, S_1|_{G_1} = t_1, \\
&\quad p(s_1, t_1), t_1 \circ t_2 |_G = t\})\}) \\
&= \text{agg}_i^2(\bigcup_{t_2 \in R_2} \{\text{agg}_i^1(\{e_i(s_1 \circ t_2) | s_1 \in R_1, s_1 \circ t_2 |_G = t, p(s_1 \circ t_2)\})\}) \\
&= \text{agg}_i^2(\bigcup_{t_2 \in R_2} \{\text{agg}_i^1(\{e_i(t_1 \circ t_2) | t_1 \in R_1, t_1 \circ t_2 |_G = t, p(t_1 \circ t_2)\})\}) \\
&= \text{agg}_i(\bigcup_{t_2 \in R_2} \{e_i(t_1 \circ t_2) | t_1 \in R_1, p(t_1 \circ t_2), t_1 \circ t_2 |_G = t\}) \\
&= \text{agg}_i(\{e_i(s) | s \in \bigcup_{t_2 \in R_2} \{t_1 \circ t_2 | t_1 \in R_1, p(t_1 \circ t_2)\}, S|_G = t\}) \\
&= a_i
\end{aligned}$$

Equivalence 20.1 can be used to add an additional coalescing grouping in front of any of a sequence of joins. Consider the schematic operator tree in Figure 20.2(a). It is equivalent to the one in (b), which in turn is equivalent to the one in (c) if the preconditions of Equivalence 20.1 hold. Performing a similar operation multiple times, any of the join operations can be made to be preceded by a coalescing grouping.

19.6 More Possibilities

Yan and Larson provide a whole set of possibilities to reorder grouping and join [825]. The query they consider is one of the form

```

select [all|distinct]  $P_1, P_2, \vec{A}_1, \vec{A}_2$ 
from  $R_1, R_2$ 
where  $p_1$  and  $p_{1,2}$  and  $p_2$ 
group by  $G_1, G_2$ 

```

The abbreviations used are explained below. Note that the order or grouping of entries in the different clauses is of no relevance. It was just introduced by us for convenience. An overview of all the different plans that can be produced from an initial plan is given in Figure ???. Every of the subsequent equivalences has the initial plan as its left-hand side and one of the other plans as its right-hand




Figure 19.2: Applications of coalescing grouping

side. Before giving the equivalences together with their conditions, we define some notation, some of them already used in the query above:

R_1 is a relation that contains aggregation columns

R_2 is a relation that may or may not contain aggregation columns

P_i are the selection columns of R_i . Define $P = P_1 \cup P_2$.

\vec{A}_1 is a possibly empty sequence of aggregation function calls on columns of R_1 . More specifically, we assume that \vec{A}_1 is of the form $a_1 : \text{agg}_1(e_1), \dots, a_k : \text{agg}_k(e_k)$ and require that $\mathcal{F}(\vec{A}_1) \subseteq \mathcal{A}(R_1)$

\vec{A}_2 is a possibly empty sequence of aggregation function calls on columns of R_2 . More specifically, we assume that \vec{A}_2 is of the form $a_{k+1} : \text{agg}_1(e_{k+1}), \dots, a_n : \text{agg}_k(e_n)$ and require that $\mathcal{F}(\vec{A}_2) \subseteq \mathcal{A}(R_2)$

\vec{A} is the concatenation of \vec{A}_1 and \vec{A}_2 . \vec{A} may be empty.

F_i are the aggregation columns of R_1 and possibly R_2 . That is, $\vec{A}^I = \mathcal{F}(\vec{A}_1)$ and $\vec{A}^O = \mathcal{F}(\vec{A}_2)$. Define $F = \vec{A}^I \cup \vec{A}^O$. Then, in case of *eager/lazy group-by*, *eager/lazy count*, and *double eager/lazy*, $F \subseteq \mathcal{A}(R_1)$. In case of *eager/lazy group-by-count* and *eager/lazy split*, \vec{A}^O may contribute to $F = \vec{A}^I \cup \vec{A}^O$.

A_i contains the columns defined by \vec{A}_i . More precisely, $\vec{A}_1 = a_1 : \text{agg}_1(e_1), \dots, a_k : \text{agg}_k(e_k)$ and $\vec{A}_2 = a_1 : \text{agg}_1(e_1), \dots, a_k : \text{agg}_k(e_k)$. resulting columns of the application of F on AA in the first group-by when *eager group-by* is performed on the above query.

p_i is a selection predicate on columns of R_i

$p_{1,2}$ is the join predicate with columns from R_1 and R_2 .

G_i are the grouping columns of R_i ($G_i \subseteq \mathcal{A}(R_i)$). Define $G = G_1 \cup G_2$.

G_i^+ are the columns of R_i participating in join and grouping, i.e. $G_i^+ := (G_i \cup \mathcal{F}(p_{1,2})) \cap \mathcal{A}(R_i)$

The query can be translated into the algebra as follows:

$$\Pi_{P,A}^{(D)}(\Gamma_{G;\vec{A}}(\sigma_{p_1 \wedge p_{1,2} \wedge p_2}(R_1 \times R_2)))$$

where the projection is duplicate eliminating if and only if the query specifies **select distinct**. P is allowed to contain more columns than those in G if these are functionally determined by G .

For the equivalences to follow, we assume that duplicates are preserved. That is, the algebraic query representation is

$$\Pi_{P,A}(\Gamma_{G;\vec{A}}(\sigma_{p_1 \wedge p_{1,2} \wedge p_2}(R_1 \times R_2)))$$

We will skip any leading Π from subsequent expressions.

19.6.1 Eager/Lazy Group-By-Count

In the following equivalence:

H_1 denotes a set of columns in R_1

c is the column produced by **count(*)** after grouping $\sigma_{p_1}(R_1)$ on H_1

A_1 the rest of the columns produced by \vec{A}_1 in the first group-by of table $\sigma_{p_1}(R_1)$ on H_1

F_{ua} the duplicated aggregation function of F_u

Further:

$$\begin{aligned} AA &= AA_d \cup AA_u \\ AA_d &= AA \cap \mathcal{A}(R_1) \\ AA_u &= AA \cap \mathcal{A}(R_2) \\ F &= A_1 \cup A_2 \end{aligned}$$

where A_1 applies to AA_d and A_2 applies to AA_u .

The expressions $E_1 :=$

$$F[AA_d, AA_u] \Pi_A[GA_d, GA_u, AA_d, AA_u] \mathcal{G}[GA_d, GA_u] \sigma[p_1 \wedge p_{1,2} \wedge p_2](R_1 \times R_2)$$

and $E_2 :=$

$$\begin{aligned} & \Pi_d[GA_d, GA_u, FAA] \\ (F_{ua}[AA_u, CNT], F_{d2}[FAA_d])\Pi_A[GA_d, GA_u, AA_u, FAA_d, CNT]\mathcal{G}[GA_d, GA_u] \\ & \sigma[p_{1,2}, p_2](E_3) \end{aligned}$$

with $E_3 :=$

$$((F_{d1}[AA_d], COUNT[])\Pi_A[H_1, GA_d^+, AA_d]\mathcal{G}[H_1]\sigma[p_1]R_1) \times R_2)$$

are equivalent if

1. A_1 contains only decomposable aggregation functions and can be decomposed into F_{d1} and F_{d2}
2. A_2 contains Class C or Class D aggregation functions
3. $H_1 \longrightarrow G_1^+$ holds in $\sigma[p_1]R_1$

Note that the equivalence assumes that duplicates are preserved (no Π_D (default in paper is Π_A which is not written) at the beginning of E_1).

The main theorem requires that the final selection columns are the same as the grouping columns, and that the final projection must not remove duplicates. This can be relaxed. The final selection columns may be a subset of the grouping columns and the final projection may remove duplicates [822].

Eager/Lazy Group-By We consider the special case where G_1 contains all the aggregation columns. In the following:

H_1 denotes a set of columns in R_1

A_1 denotes the columns produced by applying $F[AA]$ after grouping table R_1 on H_1 .

Then, the expressions $E_1 :=$

$$F[AA]\Pi_A[GA_d, GA_u, AA]\mathcal{G}[GA_d, GA_u]\sigma[p_1 \wedge p_{1,2} \wedge p_2](R_1 \times R_2)$$

and $E_2 :=$

$$\vec{A}^O[FAA_d]\Pi_A[GA_d, GA_u, FAA_d]\mathcal{G}[GA_d, GA_u]\Pi_A[GA_d, GA_u, FAA_d]\sigma[p_{1,2} \wedge p_2](E_3)$$

where $E_3 :=$

$$(\vec{A}^I[AA]\Pi_A[H_1, GA_d^+, AA]\mathcal{G}[H_1]\sigma[p_1]R_1) \times R_2$$

are equivalent if

1. all aggregation functions in $F[AA]$ are decomposable and can be decomposed into \vec{A}^I and \vec{A}^O
2. $H_1 \longrightarrow G_1^+$ holds in $\sigma[p_1]R_1$

Proof:

- Since AA_u is empty, $F_{ua}[AA_u, CNT]$ is empty
- Deleting all terms related to AA_u in E_2 of the main theorem gives the E_2 here.

Eager/Lazy Count and Eager/Lazy Distinct We consider the special case where G_2 contains all the aggregation columns. In the following equivalence:

H_1 denotes a set of grouping (sic!in YaLa95 for the first time) columns in R_1

CNT denotes the column produced by **count(*)** after grouping $\sigma[p_1]R_1$ on H_1

Then, the expressions $E_1 :=$

$$F[AA]\Pi_A[GA_d, GA_u, AA]\mathcal{G}[GA_d, GA_u]\sigma[p_1 \wedge p_{1,2} \wedge p_2](R_1 \times R_2)$$

and $E_2 :=$

$$F_a[AA, CNT]\Pi_A[GA_d, GA_u, AA, CNT] \\ \mathcal{G}[GA_d, GA_u]\Pi_A[GA_d, GA_u, AA, CNT]\sigma[p_{1,2} \wedge p_2](E_3)$$

where $E_3 :=$

$$(COUNT[]\Pi_A[H_1, GA_d^+]\mathcal{G}[H_1]\sigma[p_1]R_1) \times R_2$$

are equivalent if

1. F are Class C or Class D aggregation functions
2. $H_1 \longrightarrow G_1^+$ holds in $\sigma[p_1]R_1$

$COUNT[]$ means that we add $CNT:COUNT(*)$ to the select list of the subquery block.

If within the equivalence F only contains Class D aggregation functions, we can simply use a **distinct** in the subquery block. We then call the transformation from E_1 to E_2 *eager distinct* and its reverse application *lazy distinct*. Note that in this case F_a is the same as F .

Proof:

- Since AA_d is empty, all of A_1 , F_{d1} , and F_{d2} are empty.
- Removing all terms related to AA_d in E_2 of the main theorem gives the E_2 here.

Double Eager/Lazy In the following equivalence

H_2 denotes a set of columns in R_2

H_1 denotes a set of grouping columns in R_1

FFA denotes the columns produced by \vec{A}^I in the first group-by of table $\sigma[p_1]R_1$ on H_1

CNT denotes the column produced by **count(*)** after grouping $\sigma[p_2]R_2$ on H_2

Assume that AA belongs to R_1 . Then, the expressions $E_1 :=$

$$F[AA]\Pi_A[GA_d, GA_u, AA]\mathcal{G}[GA_d, GA_u]\sigma[p_1 \wedge p_{1,2} \wedge p_2](R_1 \times R_2)$$

and $E_2 :=$

$$F_a[\vec{A}^O[FAA], CNT]\Pi_A[GA_d, GA_u, FAA, CNT]\mathcal{G}[GA_d, GA_u]\sigma[p_{1,2}](E_3 \times E_4)$$

where $E_3 :=$

$$(COUNT[]\Pi_A[H_2, GA_u^+]\mathcal{G}[H_2]\sigma[p_2]R_2)$$

and $E_4 :=$

$$(\vec{A}^I[AA]\Pi_A[H_1, GA_d^+, AA]\mathcal{G}[H_1]\sigma[p_1]R_1)$$

are equivalent if

1. $H_2 \longrightarrow G_2^+$ holds in $\sigma[p_1]R_2$
2. $H_1 \longrightarrow G_1^+$ holds in $\sigma[p_1]R_1$
3. all aggregation functions in F are decomposable and can be decomposed into \vec{A}^I and \vec{A}^O
4. all aggregation functions in F are Class C or Class D and its duplicated aggregation function is F_a

If F contains only class D aggregation functions, we can use **distinct** in the subquery block of R_2 . In this case F_a is the same as F .

Proof:

- eager/lazy group-by, then eager/lazy count

Double Group-By Push-Down/Pull-Up

The following equivalence shows when a top group-by can be eliminated. The equivalences $E_1 :=$

$$F[AA]\Pi_A[GA_d, GA_u, AA]\mathcal{G}[GA_d, GA_u]\sigma[p_1 \wedge p_{1,2} \wedge p_2](R_1 \times R_2)$$

and $E_2 :=$

$$\Pi_A[GA_d, GA_u, FAA * CNT]\mathcal{G}[GA_d, GA_u]\sigma[p_{1,2}](E_3 \times E_4)$$

where $E_3 :=$

$$(COUNT[]\Pi_A[H_2, GA_u^+]\mathcal{G}[H_2]\sigma[p_2]R_2)$$

and $E_4 :=$

$$(F[AA]\Pi_A[H_1, GA_d^+, AA]\mathcal{G}[H_1]\sigma[p_1]R_1)$$

are equivalent if

1. $H_2 \longrightarrow G_2^+$ holds in $\sigma_{p_1}(R_2)$

2. $H_1 \longrightarrow G_1^+$ holds in $\sigma_{p_1}(R_1)$
3. all aggregation functions in \vec{A} are decomposable and can be decomposed into \vec{A}^I and \vec{A}^O
4. all aggregation functions in F are Class C or Class D and its duplicated aggregation function is F_a
5. $G_1^+ \longrightarrow H_1$ holds in $\sigma[p_1]R_1$
6. $G_2^+ \longrightarrow H_2$ holds in $\sigma[p_2]R_2$
7. (G_1, G_2) functionally determine the join columns in $\sigma[p_1 \wedge p_{1,2} \wedge p_2](R_1 \times R_2)$

Proof in [822].

Eager/Lazy Split In the following equivalence

H_1 denotes a set of columns in R_1

H_2 denotes a set of columns in R_2

c_1 denotes the column produced by **count**(*) after grouping $\sigma[p_1]R_1$ on H_1

c_2 denotes the column produced by **count**(*) after grouping $\sigma[p_2]R_2$ on H_2

A_1 denotes the columns produced by A_1 in the first aggregation of table $\sigma[p_1]R_1$ on H_1

A_2 denotes the columns produced by A_2 in the first aggregation of table $\sigma[p_2]R_2$ on H_2

F_{da} denotes the duplicated aggregation function of A_1

F_{ua} denotes the duplicated aggregation function of A_2

Assume also that

1. $AA = AA_d \cup AA_u$ where AA_d contains only columns of R_1 and AA_u contains only columns of R_2
2. $F = A_1 \cup A_2$ where A_1 applies to AA_d and A_2 applies to AA_u

Then, the expressions $E_1 :=$

$$F[AA_d, AA_u] \Pi_A[GA_d, GA_u, AA_d, AA_u] \mathcal{G}[GA_d, GA_u] \sigma[p_1 \wedge p_{1,2} \wedge p_2](R_1 \times R_2)$$

and $E_2 :=$

$$\begin{aligned} & \Pi_d[GA_d, GA_u, FAA] \\ & (F_{ua}[F_{u2}[FAA_u], CNT_1], F_{da}[F_{d2}[FAA_d], CNT_2]) \\ & \Pi_A[GA_d, GA_u, FAA_u, FAA_d, CNT_1, CNT_2] \\ & \mathcal{G}[GA_d, GA_u] \sigma[p_{1,2} \wedge p_2](E_3 \times E_4) \end{aligned}$$

where $E_3 :=$

$$(F_{d1}[AA_d], COUNT[]) \Pi_A[H_1, GA_d^+, AA_d] \mathcal{G}[H_1] \sigma[p_1] R_1$$

and $E_4 :=$

$$(F_{u1}[AA_u], COUNT[]) \Pi_A[H_2, GA_u^+, AA_u] \mathcal{G}[H_2] \sigma[p_2] R_2$$

are equivalent, if the following conditions hold:

1. A_1 contains only decomposable aggregation functions and can be decomposed into F_{d1} and F_{d2}
2. A_2 contains only decomposable aggregation functions and can be decomposed into F_{u1} and F_{u2}
3. A_2 and A_1 contain Class C or Class D aggregation functions
4. $H_2 \longrightarrow G_2^+$ holds in $\sigma[p_2] R_2$
5. $H_1 \longrightarrow G_1^+$ holds in $\sigma[p_1] R_1$

Proof:

- perform eager/lazy group-by-count on R_1 and then eager/lazy group-by-count on R_2

19.7 Translation into Our Notation

We define a special unary grouping operator Γ that more closely resembles grouping/aggregation in relational systems. For that reason assume that a_i ($1 \leq i \leq n$) are attribute names and e_i are expressions of the form $agg_i(e'_i)$ for some aggregation functions agg_r_i . Denote by \vec{A} the sequence $a_1 : e_1, \dots, a_n : e_n$. We then define

$$\Gamma_{G; \vec{A}} := \{g \circ [a_1 : v_1, \dots, a_n : v_n] \mid g \in \Pi_G^D(e), v_i = agg_r_i(G_g)\}$$

where $G_g := \{t \mid t \in e, t.G = g\}$. Then

$$\Gamma_{G; \vec{A}} := \Pi^D(\chi_{[g_1:g_1, \dots, g_k:g_k, a_1:g.a_1, \dots, a_n:g.a_n]}(\Gamma_{g; =G; [\vec{A}]}(e)))$$

Translation Table:

Translation Table:

YaLa95	we	cmd	comment
	R	rRx	used in constructs like R_i
R_d	R_1	rRa	relation R_1
R_u	R_2	rRb	relation R_2
	P	aPx	used in constructs like P_i
SGA_d	P_1	aPa	projected columns of R_1
SGA_u	P_2	aPb	projected columns of R_2
F	\vec{A}	fAx	defined as $\vec{A}_1 \circ fA2$
F_d	\vec{A}_1	fAxa	vector of aggregate functions applied to attrs of R_1
	$a_1 : \text{agg}_1(e_1), \dots, a_k : \text{agg}_k(e_k)$	fAax	form of "
F_u	\vec{A}_2	fAb	vector of aggregate functions applied to attrs of R_2
	$a_{k+1} : \text{agg}_1(e_{k+1}), \dots, a_n : \text{agg}_k(e_n)$	fAbx	form of "
AA	F	aFx	defined as $\vec{A}^T \cup aFx$
AA_d	\vec{A}^T	aFa	$\subseteq \mathcal{A}((R_1))$
AA_u	\vec{A}^O	aFb	$\subseteq \mathcal{A}((R_2))$
FAA	A	aAx	defined as $A_1 \cup A_2$
	A_1	aAa	columns containing aggregation result
	$\langle a_1, \dots, a_k \rangle$	aAax	form of "
	A_2	aAb	columns containing aggregation result
	$\langle a_{k+1}, \dots, a_n \rangle$	aAbx	form of "
	p	pPx	used in constructs like p_i .
C_d	p_1	pPa	selection predicate on relation R_1
C_u	p_2	pPb	selection predicate on relation R_1
C_0	$p_{1,2}$	pPj	join predicate for relations R_1 and R_2
	G	aGx	defined as $G_1 \cup G_2$
GA_d	G_1	aGa	grouping columns of R_1
GA_u	G_2	aGb	grouping columns of R_2
	G^+	aGpx	defined as $G_1^+ \cup G_2^+$
GA_d^+	G_1^+	aGpa	grouping columns plus join columns of R_1
GA_u^+	G_2^+	aGpb	grouping columns plus join columns of R_2
	H	aNx	set of attributes
NGA_d	H_1	aNa	
NGA_u	H_2	aNb	
CNT	c	c	

19.8 Aggregation of Non-Equi-Join Results

19.9 Bibliography

The main source of information for this section are the papers by Yan and Larson [823, 824, 825, 826]. These papers cover the material discussed in this section although in a different notation. An informal description of some of the ideas presented here can be found in a paper by Chaudhuri and Shim [129].

All of these papers somewhat discuss the the topic of introducing the optimal placement of Grouping and Aggregation in a plan generator. Chaudhuri and Shim devoted another paper to this important topic [131]. Duplicate removal is a special case of grouping with no aggregation taking place. Already very early on, Dayal, Goodman, and Katz observed that duplicate removal, can be pushed down beneficially [189]. This finding was confirmed by Pirahesh,

Hellerstein, and Hasan [586]. Gupta, Harinarayan, and Quass [326] discusses to push down duplicate elimination into a plan by counting the number of occurring duplicates. That is, they change the representation of the bag. After joins are performed, they reverse this representation change.

Pre-Aggregation: [365, 456] cardinality estimates for pre-aggregation: [366]

ToDo: [611], [634]

Aggregates: [15, 781, 130, 131, 133, 163, 164, 188] [226, 252, 313, 326, 327, 349, 421, 422] [431, 430, 527, 540, 542, 560, 690, 704, 717] [825, 847, 500]

Chapter 20

Grouping and Aggregation

20.1 Introduction

In general, join and grouping operations are not reorderable. Consider the following relations R and S

R	A	B	S	A	C
	a	5		a	7
	a	6		a	8

Joining these relations R and S results in

$R \bowtie S$	A	B	C
	a	5	7
	a	5	8
	a	6	7
	a	6	8

Applying $\Gamma_{A;count(*)}$ to R and $R \bowtie S$ yields

$\Gamma_{A;count(*)}(R)$	A	count (*)	$\Gamma_{A;count(*)}(R \bowtie S)$	A	count (*)
	a	2		a	4

Compare this to the result of $\Gamma_{A;count(*)}(R) \bowtie S$:

$\Gamma_{A;count(*)}(R) \bowtie S$	A	count (*)	C
	a	2	7
	a	2	8

Hence $\Gamma_{A;count(*)}(R) \bowtie S \neq \Gamma_{A;count(*)}(R \bowtie S)$.

Since grouping and join operations are in general not reorderable, it is important that a query language determines the order of grouping and join operators properly. In SQL, the grouping operator is applied after the join operators of a query block.

For example, given the relations schemata

Emp (eid, name, age, salary) and

Sold (sid, eid, date, product_id, price)

and the query

```

select    e.eid, sum (s.price) as amount
from      Emp e, Sold s
where     e.eid = s.eid and
           s.date between “2000-01-01” and “2000-12-31”
group by s.eid, s.name

```

results in the algebraic expression

$$\Pi_{e.eid, amount} (\Gamma_{s.eid, amount: \text{sum}(s.price)} (Emp[e] \bowtie_{e.eid=s.eid} \sigma_p (Sold[s])))$$

where p denotes

$$s.date \geq '2000-01-01' \wedge s.date \leq '2000-12-31'$$

Figure 20.1 (a) shows this plan graphically. Note that the grouping operator is executed last in the plan.

Now consider the plan where we push the grouping operator down:

$$\Pi_{e.eid, amount} (Emp[e] \bowtie_{e.eid=s.eid} (\Gamma_{s.eid, amount: \text{sum}(s.price)} (\sigma_p (Sold[s])))$$

This plan (see also Figure 20.1 (b)) is equivalent to the former plan. Moreover, if the grouping operator strongly reduces the cardinality of

$$\sigma_{s.date \geq \dots} (Sold[s])$$

because every employee sells many items, then the latter plan might become cheaper since the join inputs are smaller than in the former plan. This motivates the search for conditions under which join and grouping operators can be reordered. Several papers discuss this reorderability [129, 823, 824, 825, 826]. We will summarize their results in subsequent sections.

20.2 Lazy and eager group by

Lazy group by pulls a group operator up over a join operator [823, 824, 825, 826]. Eager group by does the opposite. This may also be called *Push-Down Grouping* and Pull-Up Grouping.

Consider the query:

```

select[all | distinct]   $\overrightarrow{A, F(B)}$ 
from                    $R, S$ 
where                    $p_R \wedge p_S \wedge p_{R,S}$ 
group by                $G$ 

```

```

      project[e.eid, amount]

group by[e.eid; amount:sum(price)]

      join[e.eid=s.eid]

Emp[e]          select[s.date between ...]

                Sold[s]
(a)

```

```

      project[e.eid, amount]

      join[e.eid = s.eid]

Emp[e]          group by[e.eid, amount:sum(price)]

                select[s.date ...]

                Sold[s]
(b)

```

Figure 20.1: Two equivalent plans

with

$$G = G_R \cup G_S, G_R \subseteq \mathcal{A}(R), G_S \subseteq \mathcal{A}(S),$$

$$\mathcal{F}(p_R) \subseteq \mathcal{A}(R), \mathcal{F}(p_S) \subseteq \mathcal{A}(S)$$

$$\mathcal{F}(p_{R,S}) \subseteq \mathcal{A}(R) \cup \mathcal{A}(S)$$

$$B \subseteq \mathcal{A}(R) \quad A = A_R \cup A_S, A_R \subseteq G_R, A_S \subseteq G_S$$

$$\alpha_R = G_R \cup \mathcal{F}(p_{R,S}) \setminus \mathcal{A}(S) \quad \kappa_R \text{ key of } R$$

$$\alpha_S = G_S \cup \mathcal{F}(p_{R,S}) \setminus \mathcal{A}(R) \quad \kappa_S \text{ key of } S$$

We are interested in the conditions under which the query can be rewritten into

select [all distinct]	A, FB
from	R', S'
where	$p_{R,S}$
with	
	$R'(\alpha_R, FB) \equiv$
	select all $\alpha_R, \vec{F}(B)$ as FB
	from R
	where p_R
	group by α_R
and	
	$S'(\alpha_S) \equiv$
	select all α_R
	from S
	where p_S

The following equivalence expresses this rewrite in algebraic terms.

$$\Pi_{A,F}^{[d]} \left(\Gamma_{G;F:\vec{F}(B)} \left(\sigma_{p_R}(R) \bowtie_{p_{R,S}} \sigma_{p_S}(S) \right) \right) \equiv$$

$$\Pi_{A,F}^{[d]} \left(\Gamma_{\alpha_R;F:\vec{F}(B)} \left(\sigma_{p_R}(R) \right) \bowtie_{p_{R,S}} \sigma_{p_S}(S) \right)$$

holds iff in $\sigma_{p_R \wedge p_S \wedge p_{R,S}}(R \times S)$ the following functional dependencies hold:

$$FD_1 : G \rightarrow \alpha_R$$

$$FD_2 : \alpha_R, G_S \rightarrow \kappa_S$$

Note that since $G_S \subseteq G$, this implies $G \rightarrow \kappa_S$.

FD_2 implies that for any group there is at most one join partner in S . Hence, each tuple in $\Gamma_{\alpha_R;F:\vec{F}(B)}(\sigma_{p_R}(R))$ contributes at most one row to the overall result.

FD_1 ensures that each group of the expression on the left-hand side corresponds to at most one group of the group expression on the right-hand side.

We now consider queries with a **having** clause.

In addition to the assumptions above, we have that the tables in the **from** clause can be partitioned into R and S such that R contains all aggregated

columns of both the **select** and the **having** clause. We further assume that conjunctive terms in the **having** clause that do not contain aggregate functions have been moved to the **where** clause.

Let the predicate of the **having** clause have the form $H_R \wedge H_0$ where $H_R \subseteq \mathcal{A}(R)$ and $H_0 \subseteq R \cup S$ where H_0 only contains non-aggregated columns from S .

We now consider all queries of the form

$$\begin{array}{ll}
 \mathbf{select}[\mathbf{all} \mid \mathbf{distinct}] & A, \vec{F}(B) \\
 \mathbf{from} & R, S \\
 \mathbf{where} & p_R \wedge p_S \wedge p_{R,S} \\
 \mathbf{group\ by} & G \\
 \mathbf{having} & H_0 \left(\vec{F}_0(B) \right) \wedge H_R \left(\vec{F}_R(B) \right)
 \end{array}$$

where \vec{F}_0 and \vec{F}_R are vectors of aggregate functions on the aggregated columns B .

An alternative way to express such a query is

$$\begin{array}{ll}
 \mathbf{select}[\mathbf{all} \mid \mathbf{distinct}] & G, FB \\
 \mathbf{from} & R', S \\
 \mathbf{where} & c_S \wedge c_{R,S} \wedge H_0(F_0B) \\
 \\
 \mathbf{where} & R'(\alpha_R, FB, F_0B) \equiv \\
 \\
 \mathbf{select\ all} & \alpha_R, \vec{F}(B) \text{ as } FB, \vec{F}_0(B) \text{ as } F_0B \\
 \mathbf{from} & R \\
 \mathbf{where} & c_R \\
 \mathbf{group\ by} & \alpha_R \\
 \mathbf{having} & H_R \left(\vec{F}_R(B) \right)
 \end{array}$$

The according equivalence is [825]:

$$\begin{aligned}
 & \Pi_{G,F} \left(\sigma_{H_R \wedge H_0} \left(\Gamma_{G;F:\vec{F}(B),F_R:\vec{F}_R(B),F_0:\vec{F}_0(B)} \left(\sigma_{p_R \wedge p_S \wedge p_{R,S}} (R \times S) \right) \right) \right) \\
 & \equiv \\
 & \Pi_{G,F} \left(\sigma_{p_{R,S} \wedge p_S \wedge H_0(F_0)} \right) \left(\Pi_{G,F,F_0} \left(\sigma_{H_R} \left(\Gamma_{G;F:\vec{F}(B),F_R:\vec{F}_R(B),F_0:\vec{F}_0(B)} (R) \right) \right) \times S \right)
 \end{aligned}$$

20.3 Coalescing Grouping

In this section we introduce *coalescing grouping* which slightly generalizes *simple coalescing grouping* as introduced in [129].

We first illustrate the main idea by means of an example.

Given two relation schemes

Sales (pid, deptid, total_price)
 Department (did, name, region)

the query

```

select    region, sum (total_price) as s
from      Sales, Department
where     did = deptid
group by  region

```

is straightforwardly translated into the following algebraic expression:

$$\Gamma_{region;s:sum(total_price)}(Sales \bowtie_{deptid=did} Department)$$

Note that Equivalence ?? cannot be applied here. However, if there are many sales performed by a department, it might be worth reducing the cardinality of the left join input by introducing an additional group operator. The result is

$$\Gamma_{region;s=sum(s')}(\Gamma_{deptid;s':sum(total_price)}(Sales) \bowtie_{deptid=did} Department)$$

Note that we must keep the outer grouping.

That is, we introduced an additional group operator to reduce the cardinality of sales. This way, all subsequent joins (only one in this case) become cheaper and the additional group operator may result in a better plan.

We have the following restrictions for this section:

1. There are no NULL-values allowed for attributes occurring in the query.
2. All queries are of the form **select all**.
That is **select distinct** is not allowed.
3. All aggregate functions *agg* must fulfill $agg_{s_1} \cup s_2 = agg\{agg(s_1), agg(s_2)\}$ for bags s_1 and s_2 .
This has two consequences:
 - Allowed are only sum, min, max. Not allowed are avg and count.
 - For any allowed aggregate function we only allow for **agg(all ...)**.
Forbidden is **agg(distinct ...)**.
4. The query is a single-block conjunctive query with no **having** and no **order by** clause.

The above transformation is an application of the following equivalence, where R_1 and R_2 can be arbitrary algebraic expressions:

$$\Gamma_{G;A}(R_1 \bowtie_p R_2) \equiv \Gamma_{G;A_2}(\Gamma_{G_1;A_1}(R_1) \bowtie_p R_2) \quad (20.1)$$

with

$$\begin{aligned}
A &= A_1 : agg_1(e_1), \dots, A_n : agg_n(e_n) \\
A_1 &= A_1^1 : agg_1^1(e_1), \dots, A_n^1 : agg_n^1(e_n) \\
A_2 &= A_1 : agg_1^2(A_1^1), \dots, A_n : agg_n^2(A_n^1) \\
G_1 &= (\mathcal{F}(p) \cup G) \cap \mathcal{A}(R_1)
\end{aligned}$$

Further, the following condition must hold for all $i(1 \leq i \leq n)$:

$$agg_i \left(\bigcup_k S_k \right) = agg_i^2 \left(\bigcup_k \{agg_i^1(S_i)\} \right)$$

In the above example, we had $agg_1 = agg_1^1 = agg_1^2 = \mathbf{sum}$.

We now prove the correctness of Equivalence 20.1.

Proof:

First, note that

$$R_1 \bowtie_p R_2 = \bigcup_{t_2 \in R_2} R_1 \bowtie_p \{t_2\} \tag{20.2}$$

Second, note that for a given t_2

$$\begin{aligned} \Gamma_{G;A}(R_1[t_1]) \bowtie_p \{t_2\} &= \sigma_{p(t_1 \circ t_2)}(\Gamma_{G;A}(R_1[t_1])) \\ &= \Gamma_{G;A}(\sigma_{p(t_1 \circ t_2)}(R_1[t_1])) \\ &= \Gamma_{G;A}(R_1[t_1] \bowtie_p \{t_2\}) \end{aligned} \tag{20.3}$$

holds where we have been a little sloppy with t_1 . Applying (20.2) and (20.3) to $\Gamma_{G_1;A_1}(R_1) \bowtie_p R_2$, the inner part of the right-hand side of the equivalence yields:

$$\begin{aligned} \Gamma_{G_1;A_1}(R_1) \bowtie_p R_2 &= \bigcup_{t_2 \in R_2} \Gamma_{G_1;A_1}(R_1) \bowtie_p \{t_2\} \\ &= \bigcup_{t_2 \in R_2} \Gamma_{G_1;A_1}(R_1 \bowtie_p \{t_2\}) \end{aligned} \tag{20.4}$$

Call the last expression X.

Then the right-hand side of our equivalence becomes

$$\begin{aligned} \Gamma_{G;A_2}(X) &= \{t \circ a_2 \mid t \in \Pi_G(X), a_2 = (A_1 : a_1^2, \dots, A_n : a_n^2), \\ &\quad a_i^2 = agg_i^2(\{s.A_i^1 \mid s \in X, S|_G = t\})\} \end{aligned} \tag{20.5}$$

Applying (20.2) to the left-hand side of the equivalence yields:

$$\Gamma_{G;A}(R_1 \bowtie_p R_2) = \Gamma_{G;A} \left(\bigcup_{t_2 \in R_2} \{t_1 \circ t_2 \mid t_1 \in R_1, p(t_1 \circ t_2)\} \right) \tag{20.6}$$

Abbreviate $\bigcup_{t_2 \in R_2} \{t_1 \circ t_2 \mid t_1 \in R_1, p(t_1 \circ t_2)\}$ by Y.

Applying the definition of $\Gamma_{G;A}$ yields:

$$\begin{aligned} \{t \circ a \mid t \in \Pi_G(Y), a = (A_1 : e_1, \dots, A_n : e_n), \\ a_i = agg_i(\{e_i(s) \mid s \in Y, S|_G = t\})\} \end{aligned} \tag{20.7}$$

Compare (20.5) and (20.7). Since $\Pi_G(X) = \Pi_G(Y)$, they can only differ in their values of A_i .

Hence, it suffices to prove that $a_i^2 = a_i$ for $1 \leq i \leq n$ for any given t .

$$\begin{aligned}
a_i^2 &= agg_i^2(\{s.A_i^1 | s \in X, S|_G = t\}) \\
&= agg_i^2(\{s.A_i^1 | s \in \bigcup_{t_2 \in R_2} \Gamma_{G_1; A_1}(R_1 \bowtie_p \{t_2\}), S|_G = t\}) \\
&= agg_i^2(\{s.A_i^1 | s \in \bigcup_{t_2 \in R_2} \{t_1 \circ t_2 \circ a_1 | t_1 \in \Pi_{G_1}(R_1), p(t_1 \circ t_2), \\
&\quad a_1 = (A_1^1 : a_1^1, \dots, A_n^1 : a_n^1) \\
&\quad a_i^1 = agg_i^1(\{e_i(s_1 \circ t_2) | s_1 \in R_1, S_1|_{G_1=t_1}, p(s_1, t_2)\}) \\
&\quad S|_G = t\}) \\
&= agg_i^2(\bigcup_{t_2 \in R_2} \{agg_i^1(\{e_i(s_1 \circ t_2) | t_1 \in \Pi_{G_1}(R_1), p(t_1 \circ t_2), s_1 \in R_1, S_1|_{G_1} = t_1, \\
&\quad p(s_1, t_1), t_1 \circ t_2|_G = t\})\}) \\
&= agg_i^2(\bigcup_{t_2 \in R_2} \{agg_i^1(\{e_i(s_1 \circ t_2) | s_1 \in R_1, s_1 \circ t_2|_G = t, p(s_1 \circ t_2)\})\}) \\
&= agg_i^2(\bigcup_{t_2 \in R_2} \{agg_i^1(\{e_i(t_1 \circ t_2) | t_1 \in R_1, t_1 \circ t_2|_G = t, p(t_1 \circ t_2)\})\}) \\
&= agg_i(\bigcup_{t_2 \in R_2} \{e_i(t_1 \circ t_2) | t_1 \in R_1, p(t_1 \circ t_2), t_1 \circ t_2|_G = t\}) \\
&= agg_i(\{e_i(s) | s \in \bigcup_{t_2 \in R_2} \{t_1 \circ t_2 | t_1 \in R_1, p(t_1 \circ t_2)\}, S|_G = t\}) \\
&= a_i
\end{aligned}$$

Equivalence 20.1 can be used to add an additional coalescing grouping in front of any of a sequence of joins. Consider the schematic operator tree in Figure 20.2(a). It is equivalent to the one in (b), which in turn is equivalent to the one in (c) if the preconditions of Equivalence 20.1 hold. Performing a similar operation multiple times, any of the join operations can be made to be preceded by a coalescing grouping.

20.4 ToDo

[611]

Γ	Γ	Γ
\bowtie	\bowtie	\bowtie
$\bowtie \cancel{a}$	$e \bowtie$	$\Gamma_c \bowtie$
$\bowtie \cancel{b}$	$\bowtie \cancel{a}$	$e \bowtie \cancel{a}$
$e \cancel{e}$	$\cancel{e} \cancel{b}$	$\cancel{e} \cancel{b}$
(a)	(b)	(c)

Figure 20.2: Applications of coalescing grouping

Chapter 21

Generating Plans for the Full Algebra

Chapter 22

Generating DAG-structured Plans

```
@misc{ roy-optimization,  
  author = "Prasan Roy",  
  title = "Optimization of DAG-Structured Query Evaluation Plans",  
  url = "citeseer.nj.nec.com/roy98optimization.html" }
```


Chapter 23

Deriving and Dealing with Interesting Orderings and Groupings

[This chapter was written by Thomas Neumann and Guido Moerkotte]

23.1 Introduction

The most expensive operations (e.g. join, grouping, duplicate elimination) during query evaluation can be performed more efficiently if the input is ordered or grouped in a certain way. Therefore, it is crucial for query optimization to recognize cases where the input of an operator satisfies the ordering or grouping requirements needed for a more efficient evaluation. Since a plan generator typically considers millions of different plans – and, hence, operators –, this recognition easily becomes a performance bottleneck for plan generation, often leading to heuristic solutions.

The importance of exploiting available orderings has already been recognized in the seminal work of Selinger et al [672]. They presented the concept of interesting orderings and showed how redundant sort operations could be avoided by reusing available orderings, rendering sort-based operators like sort-merge join much more interesting.

Along these lines, it is beneficial to reuse available grouping properties, for example for hash-based operators. While heuristic techniques to avoid redundant group-by operators have been given [129], for a long time groupings have not been treated as thoroughly as orderings. One reason might be that while orderings and groupings are related (every ordering is also a grouping), groupings behave somewhat differently. For example, a tuple stream grouped on the attributes $\{a, b\}$ need not be grouped on the attribute $\{a\}$. This is different from orderings, where a tuple stream ordered on the attributes (a, b) is also ordered on the attribute (a) . Since no simple prefix (or subset) test exists for groupings, optimizing groupings even in a heuristic way is much more difficult than optimizing orderings. Still, it is desirable to combine order optimization and the optimization of groupings, as the problems are related and treated sim-

ilarly during plan generation. Recently, some work in this direction has been published [790]. However, this only covers a special case of grouping. Instead, in this chapter we follow the approach presented by Neumann and Moerkotte [544, 543]

Other existing frameworks usually consider only order optimization, and experimental results have shown that the costs for order optimization can have a large impact on the total costs of query optimization [544]. Therefore, some care is needed when adding groupings to order optimization, as a slowdown of plan generation would be unacceptable.

In this chapter, we present a framework to efficiently reason about orderings and groupings. It can be used for the plan generator described in Chapter ??, but is actually an independent component that could be used in any kind of plan generator. Experimental results show that it efficiently handles orderings and groupings at the same time, with no additional costs during plan generation and only modest one time costs. Actually, the operations needed for both ordering and grouping optimization during plan generation can be performed in $O(1)$, basically allowing to exploit groupings for free.

23.2 Problem Definition

The order manager component used by the plan generator combines order optimization and the handling of grouping in one consistent set of algorithms and data structures. In this section, we give a more formal definition of the problem and the scope of the framework. First, we define the operations of ordering and grouping (Section 23.2.1 and 23.2.2). Then, we briefly discuss functional dependencies (Section 23.2.3) and how they interact with algebraic operators (Section 23.2.4). Finally, we explain how the component is actually used during plan generation (Section 23.2.5).

23.2.1 Ordering

During plan generation, many operators require or produce certain orderings. To avoid redundant sorting, it is required to keep track of the orderings a certain plan satisfies. The orderings that are relevant for query optimization are called *interesting orders* [672]. The set of *interesting orders* for a given query consists of

1. all orderings required by an operator of the physical algebra that may be used in a query execution plan for the given query, and
2. all orderings produced by an operator of the physical algebra that may be used in a query execution plan for the given query.

This includes the final ordering requested by the given query, if this is specified.

The interesting orders are *logical orderings*. This means that they specify a condition a tuple stream must meet to satisfy the given ordering. In contrast, the *physical ordering* of a tuple stream is the actual succession of tuples in the stream. Note that while a tuple stream has only one physical ordering,

it can satisfy multiple logical orderings. For example, the stream of tuples $((1, 1), (2, 2))$ with schema (a, b) has one physical ordering (the actual stream), but satisfies the logical orderings a, b, ab and ba .

Some operators, like `sort`, actually influence the physical ordering of a tuple stream. Others, like `select`, only influence the logical ordering. For example, a `sort[a]` produces a tuple stream satisfying the ordering (a) by actually changing the physical order of tuples. After applying `select[a=b]` to this tuple stream, the result satisfies the logical orderings $(a), (b), (a, b), (b, a)$, although the physical ordering did not change. Deduction of logical orderings can be described by using the well-known notion of *functional dependency* (FD) [709]. In general, the influence of a given algebraic operator on a set of logical orderings can be described by a set of functional dependencies.

We now formalize the problem. Let $R = (t_1, \dots, t_r)$ be a stream (ordered sequence) of tuples in attributes A_1, \dots, A_n . Then R satisfies the logical ordering $o = (A_{o_1}, \dots, A_{o_m})$ ($1 \leq o_i \leq n$) if and only if for all $1 \leq i < j \leq r$ the following condition holds:

$$\begin{aligned} & (t_i.A_{o_1} \leq t_j.A_{o_1}) \\ \wedge \quad & \forall 1 < k \leq m \quad (\exists 1 \leq l < k (t_i.A_{o_l} < t_j.A_{o_l})) \vee \\ & ((t_i.A_{o_{k-1}} = t_j.A_{o_{k-1}}) \wedge \\ & (t_i.A_{o_k} \leq t_j.A_{o_k})) \end{aligned}$$

Next, we need to define the inference mechanism. Given a logical ordering $o = (A_{o_1}, \dots, A_{o_m})$ of a tuple stream R , then R obviously satisfies any logical ordering that is a prefix of o including o itself.

Let R be a tuple stream satisfying both the logical ordering $o = (A_1, \dots, A_n)$ and the functional dependency $f = B_1, \dots, B_k \rightarrow B_{k+1}$ ¹ with $B_i \in \{A_1 \dots A_n\}$. Then R also satisfies any logical ordering derived from o as follows: add B_{k+1} to o at any position such that all of B_1, \dots, B_k occurred before this position in o . For example, consider a tuple stream satisfying the ordering (a, b) ; after inducing the functional dependency $a, b \rightarrow c$, the tuple stream also satisfies the ordering (a, b, c) , but not the ordering (a, c, b) . Let O' be the set of all logical orderings that can be constructed this way from o and f after prefix closure. Then, we use the following notation: $o \vdash_f O'$. Let e be the equation $A_i = A_j$. Then, $o \vdash_e O'$, where O' is the prefix closure of the union of the following three sets. The first set is O_1 defined as $o \vdash_{A_i \rightarrow A_j} O_1$, the second is O_2 defined as $o \vdash_{A_j \rightarrow A_i} O_2$, and the third is the set of logical orderings derived from o where a possible occurrence of A_i is replaced by A_j or vice versa. For example, consider a tuple stream satisfying the ordering (a) ; after inducing the equation $a = b$, the tuple stream also satisfies the orderings $(a, b), (b)$ and (b, a) . Let e be an equation of the form $A = \text{const}$. Then O' ($o \vdash_e O'$) is derived from o by inserting A at any position in o . This is equivalent to $o \vdash_{\emptyset \rightarrow A} O'$. For example, consider a tuple stream satisfying the ordering (a, b) ; after inducing the equation $c = \text{const}$ the tuple stream also satisfies the orderings $(c, a, b), (a, c, b)$ and (a, b, c) .

¹Any functional dependency which is not in this form can be normalized into a set of FDs of this form.

Let O be a set of logical orderings and F a set of functional dependencies (and possibly equations). We define the sets of inferred logical orderings $\Omega_i(O, F)$ as follows:

$$\begin{aligned}\Omega_0(O, F) &:= O \\ \Omega_i(O, F) &:= \Omega_{i-1}(O, F) \cup \\ &\quad \bigcup_{f \in F, o \in \Omega_{i-1}(O, F)} O' \text{ with } o \vdash_f O'\end{aligned}$$

Let $\Omega(O, F)$ be the prefix closure of $\bigcup_{i=0}^{\infty} \Omega_i(O, F)$. We write $o \vdash_F o'$ if and only if $o' \in \Omega(O, F)$.

23.2.2 Grouping

It was shown in [790] that, similar to order optimization, it is beneficial to keep track of the groupings satisfied by a certain plan. Traditionally, group-by operators are either applied after the rest of the query has been processed or are scheduled using some heuristics [129]. However, the plan generator could take advantage of grouping properties produced e.g. by avoiding re-hashing if such information was easily available.

Analogous to order optimization, we call this *grouping optimization* and define that the set of *interesting groupings* for a given query consists of

1. all groupings required by an operator of the physical algebra that may be used in a query execution plan for the given query
2. all groupings produced by an operator of the physical algebra that may be used in a query execution plan for the given query.

This includes the grouping specified by the group-by clause of the query, if any exists.

These groupings are similar to logical orderings, as they specify a condition a tuple stream must meet to satisfy a given grouping. Likewise, functional dependencies can be used to infer new groupings.

More formally, a tuple stream $R = (t_1, \dots, t_r)$ in attributes A_1, \dots, A_n satisfies the grouping $g = \{A_{g_1}, \dots, A_{g_m}\}$ ($1 \leq g_i \leq n$) if and only if for all $1 \leq i < j < k \leq r$ the following condition holds:

$$\begin{aligned}\forall 1 \leq l \leq m \quad t_i.A_{g_l} &= t_k.A_{g_l} \\ \Rightarrow \forall 1 \leq l \leq m \quad t_i.A_{g_l} &= t_j.A_{g_l}\end{aligned}$$

Two remarks are in order here. First, note that a grouping is a set of attributes and not – as orderings – a sequence of attributes. Second, note that given two groupings g and $g' \subset g$ and a tuple stream R satisfying the grouping g , R need not satisfy the grouping g' . For example, the tuple stream $((1, 2), (2, 3), (1, 4))$ with the schema (a, b) is grouped by $\{a, b\}$, but not by $\{a\}$. This is different from orderings, where a tuple stream satisfying an ordering o also satisfies all orderings that are a prefix of o .

New groupings can be inferred by functional dependencies as follows: Let R be a tuple stream satisfying both the grouping $g = \{A_1, \dots, A_n\}$ and the functional dependency $f = B_1, \dots, B_k \rightarrow B_{k+1}$ with $\{B_1, \dots, B_k\} \subseteq \{A_1, \dots, A_n\}$. Then R also satisfies the grouping $g' = \{A_1, \dots, A_n\} \cup \{B_{k+1}\}$. Let G' be the set of all groupings that can be constructed this way from g and f . Then we use the following notation: $g \vdash_f G'$. For example $\{a, b\} \vdash_{a,b \rightarrow c} \{a, b, c\}$. Let e be the equation $A_i = A_j$. Then $g \vdash_e G'$ where G' is the union of the following three sets. The first set is G_1 defined as $g \vdash_{A_i \rightarrow A_j} G_1$, the second is G_2 defined as $g \vdash_{A_j \rightarrow A_i} G_2$, and the third is the set of groupings derived from g where a possible occurrence of A_i is replaced by A_j or vice versa. For example, $\{a, b\} \vdash_{b=c} \{a, c\}$. Let e be an equation of the form $A = \text{const}$. Then $g \vdash_e G'$ is defined as $g \vdash_{\emptyset \rightarrow A} G'$. For example, $\{a, b\} \vdash_{c=\text{const}} \{a, b, c\}$.

Let G be a set of groupings and F be a set of functional dependencies (and possibly equations). We define the set of inferred groupings $\Omega_i(G, F)$ as follows:

$$\begin{aligned} \Omega_0(G, F) &:= G \\ \Omega_i(G, F) &:= \Omega_{i-1}(G, F) \cup \\ &\quad \bigcup_{f \in F, g \in \Omega_{i-1}(G, F)} G' \text{ with } g \vdash_f G' \end{aligned}$$

Let $\Omega(G, F)$ be $\bigcup_{i=0}^{\infty} \Omega_i(G, F)$. We write $g \vdash_F g'$ if and only if $g' \in \Omega(G, F)$.

23.2.3 Functional Dependencies

The reasoning about orderings and groupings assumes that the set of functional dependencies is known. The process of gathering the relevant functional dependencies is described in detail in [709, 710]. Predominantly, there are three sources of functional dependencies:

1. key constraints
2. join predicates [references constraints]
3. filter predicates
4. simple expressions

However, the algorithm makes no assumption about the functional dependencies. If for some reason an operator induces another kind of functional dependency (e.g., when using TID-based optimizations [506]), this can be handled the same way. The only important fact is that we provide the set of functional dependencies as input to the algorithm.

23.2.4 Algebraic Operators

To illustrate the propagation of orderings and groupings during query optimization, we give some rules for concrete (physical) operators in Figure 23.1. As a

operator	requires	produces
$\text{scan}(R)$	-	$O(R)$
$\text{indexscan}(Idx)$	-	$O(Idx)$
$\text{map}(S, a = f(b))$	-	$\Omega(O(S), b \rightarrow a)$
$\text{select}(S, a = b)$	-	$\Omega(O(S), a = b)$
$\text{bnl-join}(S_1, S_2)$	-	$O(S_1)$
$\text{indexnl-join}(S_1, S_2)$	-	$O(S_1)$
$\text{djoin}(S_1, S_2)$	-	$O(S_1)$
$\text{sort}(S, a_1, \dots, a_n)$	-	(a_1, \dots, a_n)
$\text{group-by}(S, a_1, \dots, a_n)$	-	$\{a_1, \dots, a_n\}$
$\text{hash}(S, a_1, \dots, a_n)$	-	$\{a_1, \dots, a_n\}$
$\text{sort-merge}(S_1, S_2, \vec{a} = \vec{b})$	$\vec{a} \in O(S_1) \wedge \vec{b} \in O(S_2)$	$\Omega(O(S_1), \vec{a} = \vec{b})$
$\text{hash-join}(S_1, S_2, \vec{a} = \vec{b})$	$\vec{a} \downarrow \in O(S_1) \wedge \vec{b} \downarrow \in O(S_2)$	$\Omega(O(S_1), \vec{a} = \vec{b})$

Figure 23.1: Propagation of orderings and groupings

shorthand, we use the following notation:

$O(R)$ set of logical orderings and groupings satisfied by the physical ordering of the relation R

$O(S)$ inferred set of logical orderings and groupings satisfied by the tuple stream S

$x \downarrow$ $\{y | y \in x\}$

Note that these rules somewhat depend on the actual implementation of the operators, e.g. a blockwise nested loop join might actually destroy the ordering if the blocks are stored in hash tables. The rules are also simplified: For example, a group-by will probably compute some aggregate functions, inducing new functional dependencies. Furthermore, additional information can be derived from schema information: If the right-hand side of a dependent join (index nested loop joins are similar) produces at most one tuple, and the left-hand side is grouped on the free attributes of the right-hand side (e.g. if they do not contain duplicates) the output is also grouped on the attributes of the right-hand side. This situation is common, especially for index nested loop joins, and is detected automatically if the corresponding functional dependencies are considered. Therefore, it is important that all operators consider all functional dependencies they induce.

EXC

23.2.5 Plan Generation

To exploit available logical orderings and groupings, the plan generator needs access to the combined order optimization and grouping component, which we describe as an *abstract data type* (ADT). An instance of this abstract data type `OrderingGrouping` represents a set of logical orderings and groupings, and wherever necessary, an instance is embedded into a plan node. The main operations the abstract data type `OrderingGrouping` must provide are

1. a constructor for a given logical ordering or grouping,

2. a membership test (called `containsOrdering(LogicalOrdering)`) which tests whether the set contains the logical ordering given as parameter,
3. a membership test (called `containsGrouping(Grouping)`) which tests whether the set contains the grouping given as parameter, and
4. an inference operation (called `infer(set<FD>)`). Given a set of functional dependencies and equations, it computes a new set of logical orderings and groupings a tuple stream satisfies.

These operations can be implemented by using the formalism described before: `containsOrdering` tests for $o \in O$, `containsGrouping` tests for $o \in G$ and `infer(F)` calculates $\Omega(O, F)$ respectively $\Omega(G, F)$. Note that the intuitive approach to explicitly maintain the set of all logical orderings and groupings is not useful in practice. For example, if a sort operator sorts a tuple stream on (a, b) , the result is compatible with logical orderings $\{(a, b), (a)\}$. After a selection operator with selection predicate $x = \text{const}$ is applied, the set of logical orderings changes to $\{(x, a, b), (a, x, b), (a, b, x), (x, a), (a, x), (x)\}$. Since the size of the set increases quadratically with every additional selection predicate of the form $v = \text{const}$, a naive representation as a set of logical orderings is problematic. This led Simmen et al. to introduce a more concise representation, which is discussed in the related work section. Note that Simmen's technique is not easily applicable to groupings, and no algorithm was proposed to efficiently maintain the set of available groupings. The order optimization component described here closes this gap by supporting both orderings and groupings. The problem of quadratic growth is avoided by only implicitly representing the set.

23.3 Overview

As we have seen, explicit maintenance of the set of logical orderings and groupings can be very expensive. However, the ADT `OrderingGrouping` required for plan generation does not need to offer access to this set: It only allows to test if a given interesting order or grouping is in the set and changes the set according to new functional dependencies. Hence, it is *not* required to explicitly represent this set; an implicit representation is sufficient as long as the ADT operations can be implemented atop of it. In other words, we need not be able to reconstruct the set of logical orderings and groupings from the state of the ADT. This gives us room for optimizations.

The initial idea (see [544]) was to represent sets of logical orderings as *states* of a *finite state machine* (FSM). Roughly, a state of the FSM represents a current physical ordering and the set of logical orderings that can be inferred from it given a set of functional dependencies. The edges (transitions) in the FSM are labeled by sets of functional dependencies. They lead from one state to another, if the target state of the edge represents the set of logical orderings that can be derived from the orderings the edge's source node represents by applying the set of functional dependencies the edge is labeled with. We have




Figure 23.2: Possible FSM for orderings

to use sets of functional dependencies, since a single algebraic operator may introduce more than one functional dependency.

Let us illustrate the idea by a simple example and then discuss some problems. In Figure 23.2, an FSM for the interesting order (a, b, c) and its prefixes (remember that we need prefix closure) and the set of functional dependencies $\{b \rightarrow d\}$ is given. When a physical ordering satisfies (a, b, c) , it also satisfies its prefixes (a, b) and (a) . This is indicated by the ϵ transitions. The functional dependency $b \rightarrow d$ allows to derive the *logical* orderings (a, b, c, d) and (a, b, d, c) . This is handled by assuming that the *physical* ordering changes to either (a, b, c, d) or (a, b, d, c) . Hence, these states have to be added to the FSM. We further add the transitions induced by $\{b \rightarrow d\}$. Note that the resulting FSM is a *non-deterministic finite state machine* (NFSM).

Assume we have an NFSM as above. Then (while ignoring groupings) the state of the ADT is a state of the NFSM and the operations of the ADT can easily be mapped to the FSM. Testing for a logical ordering can be performed by checking if the node with the ordering is reachable from the current state by following ϵ edges. If the set must be changed because of a functional dependency the state is changed by following the edge labeled with the functional dependency. Of course, the non-determinism is in our way.

While remembering only the active state of the NFSM avoids the problem of maintaining a set of orderings, the NFSM is not really useful from a practical point of view, since the transitions are non-deterministic. Nevertheless, the NFSM can be considered as a special *non-deterministic finite automaton* (NFA), which consumes the functional dependencies and "recognizes" the possible physical orderings. Further, an NFA can be converted into a *deterministic finite automaton* (DFA), which can be handled efficiently. Remember that the construction is based on the power set of the NFA's states. That is, the states of the DFA are sets of states of the NFA [476]. We do not take the deviation over the finite automaton but instead lift the construction of deterministic finite automata from non-deterministic ones to finite state machines. Since this is not a traditional conversion, we give a proof of this step in Section ??.

Yet another problem is that the conversion from an NFSM to a *deterministic FSM* (DFSM) can be expensive for large NFSMs. Therefore, reducing the size of the NFSM is another problem we look at. We introduce techniques for reducing the set of functional dependencies that have to be considered and




Figure 23.3: Possible FSM for groupings

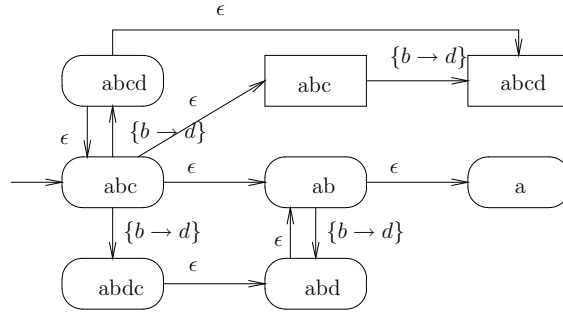


Figure 23.4: Combined FSM for orderings and groupings

further techniques to prune the NFSM in Section 23.4.7.

The idea of a combined framework for orderings and groupings was presented in [543]. Here, the main point is to construct a similar FSM for groupings and integrate it into the FSM for orderings, thus handling orderings and groupings at the same time. An example of this is shown in Figure 23.3. Here, the FSM for the grouping $\{a, b, c\}$ and the functional dependency $b \rightarrow c$ is shown. We represent states for orderings as rounded boxes and states for groupings as rectangles. Note that although the FSM for groupings has a start node similar to the FSM for orderings, it is much smaller. This is due to the fact that groupings are only compatible with themselves, no nodes for prefixes are required. However, the FSM is still non-deterministic: given the functional dependency $b \rightarrow c$, the grouping $\{a, b, c, d\}$ is compatible with $\{a, b, c, d\}$ itself and with $\{a, b, c\}$; therefore, there exists an (implicit) edge from each grouping to itself.

The FSM for groupings is integrated into the FSM for orderings by adding ϵ edges from each ordering to the grouping with the same attributes; this is due to the fact that every ordering is also a grouping. Note that although the ordering (a, b, c, d) also implies the grouping $\{a, b, c\}$, no edge is required for this, since there exists an ϵ edge to (a, b, c) and from there to $\{a, b, c\}$.

After constructing a combined FSM as described above, the full ADT supporting both orderings and groupings can easily be mapped to the FSM: The state of the ADT is a state of the FSM and testing for a logical ordering or




Figure 23.5: Possible DFSM for Figure 23.4

grouping can be performed by checking if the node with the ordering or grouping is reachable from the current state by following ϵ edges (as we will see, this can be precomputed to yield the $O(1)$ time bound for the ADT operations). If the state of the ADT must be changed because of functional dependencies, the state in the FSM is changed by following the edge labeled with the functional dependency.

However, the non-determinism of this transition is a problem. Therefore, for practical purposes the NFSM must be converted into a DFSM. The resulting DFSM is shown in Figure 23.5. Note that although in this simple example the DFSM is very small, the conversion could lead to exponential growth. Therefore, additional pruning techniques for groupings are presented in Section 23.4.7. However, the inclusion of groupings is not critical for the conversion, as the grouping part of the NFSM is nearly independent of the ordering part. In Section 23.5 we look at the size increase due to groupings. The memory consumption usually increases by a factor of two, which is the minimum expected increase, since every ordering is a grouping.

Some operators, like *sort*, change the physical ordering. In the NFSM, this is handled by changing the state to the node corresponding to the new physical ordering. Implied by its construction, in the DFSM this new physical ordering typically occurs in several nodes. For example, (a, b, c) occurs in both nodes of the DFSM in Figure 23.5. It is, therefore, not obvious which node to choose. We will take care of this problem during the construction of the NFSM (see Section 23.4.3).

23.4 Detailed Algorithm

23.4.1 Overview

Our approach consists of two phases. The first phase is the preparation step taking place before the actual plan generation starts. The output of this phase are the precomputed values used to implement the ADT. Then the ADT is used during the second phase where the actual plan generation takes place. The first phase is performed exactly once and is quite involved. Most of this section covers the first phase. Only Section 23.4.6 deals with the ADT implementation.

Figure 23.6 gives an overview of the preparation phase. It is divided into four major steps, which are discussed in the following subsections. Subsection 23.4.2 briefly reviews how the input to the first phase is determined and, more importantly, what it looks like. Section 23.4.3 describes in detail the construction of the NFSM from the input. The conversion from the NFSM to the DFSM is only briefly sketched in Section 23.4.4, for details see [476]. From the DFSM some values are precomputed which are then used for the efficient implementation of the ADT. The precomputation is described in Section 23.4.5, while their utilization and the ADT implementation are the topic of Section 23.4.6. Section 23.4.7 contains some important techniques to reduce the size of the NFSM. They are applied in Steps 2 (b), 2 (c) and 2 (e). During the discussion, we illustrate the different steps by a simple running example. More complex examples can be found in Section 23.5.

1. Determine the input
 - (a) Determine interesting orders
 - (b) Determine interesting groupings
 - (c) Determine set of functional dependencies
2. Construct the NFSM
 - (a) Construct states of the NFSM
 - (b) Filter functional dependencies
 - (c) Build filters for orderings and groupings
 - (d) Add edges to the NFSM
 - (e) Prune the NFSM
 - (f) Add artificial start state and edges
3. Convert the NFSM into a DFSM
4. Precompute values
 - (a) Precompute the compatibility matrix
 - (b) Precompute the transition table

Figure 23.6: Preparation steps of the algorithm

23.4.2 Determining the Input

Since the preparation step is performed immediately before plan generation, it is assumed that the query optimizer has already determined which indices are applicable and which algebraic operators can possibly be used to construct the query execution plan.

Before constructing the NFSM, the set of interesting orders, the set of interesting groupings and the sets of functional dependencies for each algebraic operator are determined. We denote the set of sets of functional dependencies by \mathcal{F} . It is important for the correctness of our algorithms that we note which of the interesting orders are (1) produced by some algebraic operator or (2) only tested for. Note that the interesting orders which satisfy (1) may additionally be tested for as well. We denote those orderings under (1) by O_P , those under (2) by O_T . The total set of interesting orders is defined as $O_I = O_P \cup O_T$. The orders produced are treated slightly differently in the following steps. The groupings are classified similarly to the orderings: We denote the grouping produced by some algebraic operator by G_P , and those just tested for by G_T . The total set of interesting groupings is defined as $G_I = G_P \cup G_T$. More information on how to extract interesting groupings can be found in [790]. Furthermore, for a sample query the extraction of both interesting orders and groupings is illustrated in Section 23.5.

To illustrate subsequent steps, we assume that the set of sets of functional

ToDo: details on determining interesting orders?

dependencies

$$\mathcal{F} = \{\{b \rightarrow c\}, \{b \rightarrow d\}\},$$

the interesting groupings

$$G_I = \{\{b\}\} \cup \{\{b, c\}\}$$

and the interesting orders

$$O_I = \{(b), (a, b)\} \cup \{(a, b, c)\}$$

have been extracted from the query. We assume that those in $O_T = \{(a, b, c)\}$ and $G_T = \{\{b, c\}\}$ are tested for but not produced by any operator, whereas those in $O_P = \{(b), (a, b)\}$ and $G_P = \{\{b\}\}$ may be produced by some algebraic operators.

23.4.3 Constructing the NFSM

An NFSM consists of a tuple (Σ, Q, D, q_0) , where

- Σ is the input alphabet,
- Q is the set of possible states,
- $D \subseteq Q \times (\Sigma \cup \{\epsilon\}) \times Q$ is the transition relation, and
- q_0 is the initial state.

Coarsely, Σ consists of the functional dependencies, Q of the relevant orderings and groupings, and D describes how the orderings or groupings change under a given functional dependency. Some refinements are needed to provide efficient ADT operations. The details of the construction are described now.

For the order optimization part the states are partitioned in $Q = Q_I \cup Q_A \cup \{q_0\}$, where q_0 is an artificial state to initialize the ADT, Q_I is the set of states corresponding to interesting orderings and Q_A is a set of artificial states only required for the algorithm itself. Q_A is described later. Furthermore, the set Q_I is partitioned in Q_I^P and Q_I^T , representing the orderings in O_P and O_T , respectively. To support groupings, we add to Q_I^P states corresponding to the groupings in G_P and to Q_I^T states corresponding to the groupings in G_T .

The initial NFSM contains the states Q_I of interesting groupings and orderings. For the example, this initial construction not including the start state q_0 is shown in Figure 23.7. The states representing groupings are drawn as rectangles and the states representing orderings are drawn with rounded corners.

When considering functional dependencies, additional groupings and orderings can occur. These are not directly relevant for the query, but have to be represented by states to handle transitive changes. Since they have no direct connection to the query, these states are called artificial states. Starting with the initial states Q_I , artificial states are constructed by considering functional dependencies

$$Q_A = (\Omega(O_I, \mathcal{F}) \setminus O_I) \cup (\Omega(G_I, \mathcal{F}) \setminus G_I).$$




Figure 23.7: Initial NFSM for sample query

In our example, this creates the states (b, c) and (a) , as (b, c) can be inferred from (b) when considering $\{b \rightarrow c\}$ and (a) can be inferred from (a, b) , since (a) is a prefix of (a, b) . The result is shown in Figure 23.8 (ignore the edges).

Sometimes the ADT has to be explicitly initialized with a certain ordering or grouping (e.g. after a `sort`). To support this, artificial edges are added later on. These point to the requested ordering or grouping (states in Q_I^P) and are labeled with the state that they lead to. Therefore, the input alphabet Σ consists of the sets of functional dependencies and produced orderings and groupings:

$$\Sigma = \mathcal{F} \cup Q_I^P \cup \{\epsilon\}.$$

In our example, $\Sigma = \{\{b \rightarrow c\}, \{b \rightarrow d\}, (b), (a, b), \{b\}\}$.

Accordingly, the domain of the transition relation D is

$$D \subseteq \begin{aligned} & ((Q \setminus \{q_0\}) \times (\mathcal{F} \cup \{\epsilon\}) \times (Q \setminus \{q_0\})) \\ & \cup (\{q_0\} \times Q_I^P \times Q_I^P). \end{aligned}$$

The edges are formed by the functional dependencies and the artificial edges. Furthermore, ϵ edges exist between orderings and the corresponding groupings, as orderings are a special case of grouping:

$$\begin{aligned} D_{FD} &= \{(q, f, q') \mid q \in Q, f \in \mathcal{F} \cup \{\epsilon\}, q' \in Q, q \vdash_f q'\} \\ D_A &= \{(q_0, q, q) \mid q \in Q_I^P\} \\ D_{OG} &= \{(o, \epsilon, g) \mid o \in \Omega(O_I, \mathcal{F}), g \in \Omega(G_I, \mathcal{F}), o \equiv g\} \\ D &= D_{FD} \cup D_A \cup D_{OG} \end{aligned}$$

First, the edges corresponding to functional dependencies are added (D_{FD}). In our example, this results in the NFSM shown in Figure 23.8.

Note that the functional dependency $b \rightarrow d$ has been pruned, since d does not occur in any interesting order or grouping. The NFSM can be further simplified by pruning the artificial state (b, c) , which cannot lead to a new interesting order. The result is shown in Figure 23.9. A detailed description of these pruning techniques can be found in Section 23.4.7.

The artificial start state q_0 has emanating edges incident to all states representing interesting orders in O_I^P and interesting groupings in G_I^P (D_A). Also, the states representing orderings have edges to their corresponding grouping

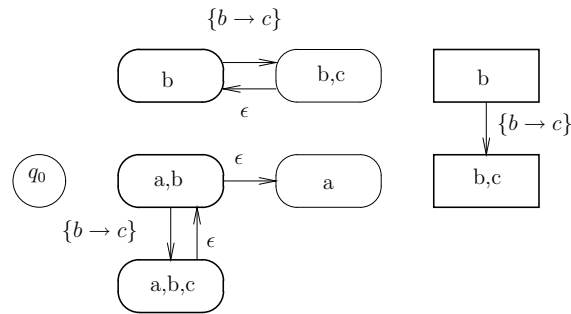


Figure 23.8: NFSM after adding D_{FD} edges




Figure 23.9: NFSM after pruning artificial states

states (D_{OG}), as every ordering is also a grouping. The final NFSM for the example is shown in Figure 23.10. Note that the states representing (a, b, c) and $\{b, c\}$ are not linked by an artificial edge since it is only tested for, as they are in Q_I^T .

23.4.4 Constructing the DFSM

The construction of the DFSM from the NFSM follows the standard power set construction that is used to translate an NFA into a DFA [476]. A formal




Figure 23.10: Final NFSM




Figure 23.11: Resulting DFSA

state	1: (a)	2: (a,b)	3: (a,b,c)	4: (b)	5: {b}	6: {b,c}
1	0	0	0	0	1	0
2	0	0	0	1	1	0
3	1	1	0	0	0	0
4	0	0	0	0	1	1
5	0	0	0	1	1	1
6	1	1	1	0	0	0

Figure 23.12: contains Matrix

description and a proof of correctness is given in Section ???. It is important to note that this construction preserves the start state and the artificial edges. The resulting DFSA for the example is shown in Figure 23.11.

23.4.5 Precomputing Values

To allow for an efficient precomputation of values, every occurrence of an interesting order, interesting grouping or set of functional dependencies is replaced by integers. This allows comparisons in constant time (equivalent entries are mapped to same integer). Further, the DFSA is represented by an adjacency matrix.

state	1: {b -> c}	2: (a, b)	3: (b)	4: {b}
q0	-	3	2	1
1	4	-	-	-
2	5	-	-	-
3	6	-	-	-
4	4	-	-	-
5	5	-	-	-
6	6	-	-	-

Figure 23.13: transition Matrix

The precomputation step itself computes two matrices. The first matrix denotes whether an NFSM state in Q_I is active, i.e. an interesting order or an interesting grouping, is contained in a specific DFSM state. This matrix can be represented as a compact bit vector, allowing tests in $O(1)$. For our running example, it is given (in a more readable form) in Figure 23.12. The second matrix contains the transition table for the DFSM relation D . Using it, edges in the DFSM can be followed in $O(1)$. For the example, the transition matrix is given in Figure 23.13.

23.4.6 During Plan Generation

During plan generation, larger plans are constructed by adding algebraic operators to existing (sub-)plans. Each subplan contains the available orderings and groupings in the form of the corresponding DFSM state. Hence, the state of the DFSM, a simple integer, is the state of our ADT `OrderingGrouping`.

When applying an operator to subplans, the ordering and grouping requirements are tested by checking whether the DFSM state of the subplan contains the required ordering or grouping of the operator. This is done by a simple lookup in the *contains* matrix.

If the operator introduces a new set of functional dependencies, the new state of the ADT is computed by following the according edge in the DFSM. This is performed by a quick lookup in the *transition* matrix.

For “atomic” subplans like table or index scans, the ordering and grouping is determined explicitly by the operator. The state of the DFSM is determined by a lookup in the transition matrix with start state q_o and the edge annotated by the produced ordering or grouping. For sort and group-by operators the state of the DFSM is determined as before by following the artificial edge for the produced ordering or grouping and then reapplying the set of functional dependencies that currently hold.

In the example, a sort on (b) results in a subplan with ordering/grouping state 2 (the state 2 is active in the DFSM), which satisfies the ordering (b) and the grouping $\{b\}$. After applying an operator which induces $b \rightarrow c$, the ordering/grouping changes to state 5 which also satisfies $\{b, c\}$.

23.4.7 Reducing the Size of the NFSM

Reducing the size of the NFSM is important for two reasons: First, it reduces the amount of work needed during the preparation step, especially the conversion from NFSM to DFSM. Even more important is that a reduced NFSM results in a smaller DFSM. This is crucial for plan generation, since it reduces the search space: Plans can only be compared and pruned if they have comparable ordering and a comparable set of functional dependencies (see [709, 710] for details). Reducing the size of the DFSM removes information that is not relevant for plan generation and, therefore, allows a more aggressive pruning of plans.

At first, the functional dependencies are pruned. Here, functional dependencies which can never lead to a new interesting order or grouping are removed.

For convenience, we extend the definition of $\Omega(O, F)$ and define

$$\Omega(O, \epsilon) := \Omega(O, \emptyset).$$

Then the set of prunable functional dependencies F_P can be described by

$$\begin{aligned} \Omega_N(o, f) &:= \Omega(\{o\}, \{f\}) \setminus \Omega(\{o\}, \epsilon) \\ F_P &:= \{f \mid f \in F \wedge \forall o \in O_I \cup G_I : \\ &\quad (\Omega(\Omega_N(o, f), F) \setminus \Omega(\{o\}, \epsilon)) \cap (O_I \cup G_I) = \emptyset\}. \end{aligned}$$

Pruning functional dependencies is especially useful, since it also prunes artificial states that would be created because of the dependencies. In the example, this removed the functional dependency $b \rightarrow d$, since d does not appear in any interesting order or grouping. This step also removes the artificial states containing d .

The artificial states are required to build the NFSM, but they are not visible outside the NFSM. Therefore, they can be pruned and merged without affecting plan generation. Two heuristics are used to reduce the set of artificial states:

1. All artificial nodes that behave exactly the same (that is, their edges lead to the same states given the same input) are merged and
2. all edges to artificial states that can reach states in Q_I only through ϵ edges are replaced with corresponding edges to the states in Q_I .

More formally, the following pairs of states can be merged:

$$\begin{aligned} \{(o_1, o_2) \mid o_1 \in Q_A, o_2 \in Q_A \wedge \forall f \in F : \\ &\quad (\Omega(\{o_1\}, \{f\}) \setminus \Omega(\{o_1\}, \epsilon)) = \\ &\quad (\Omega(\{o_2\}, \{f\}) \setminus \Omega(\{o_2\}, \epsilon))\}. \end{aligned}$$

The following states can be replaced with the next state reachable by an ϵ edge:

$$\begin{aligned} \{o \mid o \in Q_A \wedge \forall f \in F : \\ &\quad \Omega(\Omega(\{o\}, \epsilon), \{f\}) \setminus \{o\} = \\ &\quad \Omega(\Omega(\{o\}, \epsilon) \setminus \{o\}, \{f\})\}. \end{aligned}$$

In the example, this removed the state (b, c) , which was artificial and only led to the state (b) .

These techniques reduce the size of the NFSM, but still most states are artificial states, i.e. they are only created because they can be reached by considering functional dependencies when a certain ordering or grouping is available. But many of these states are not relevant for the actual query processing. For example, given a set of interesting orders which consists only of a single ordering (a) and a set of functional dependencies which consists only of $a \rightarrow b$, the NFSM will contain (among others) two states: (a) and (a, b) . The state (a, b) is created since it can be reached from (a) by considering the functional dependency, however, it is irrelevant for the plan generation, since (a, b) is not

an interesting order and is never created nor tested for. Actually, in the example above, the whole functional dependency would be pruned (since b never occurs in an interesting order), but the problem remains for combinations of interesting orders: Given the interesting orders (a) , (b) and (c) and the functional dependencies $\{a \rightarrow b, b \rightarrow a, b \rightarrow c, c \rightarrow b\}$, the NFSM will contain states for all permutations of a , b and c . But these states are completely useless, since all interesting orders consist only of a single attribute and, therefore, only the first entry of an ordering is ever tested.

Ideally, the NFSM should only contain states which are relevant for the query; since this is difficult to ensure, a heuristic can be used which greatly reduces the size of the NFSM and still guarantees that all relevant states are available: When considering a functional dependency of the form $a \rightarrow b$ and an ordering o_1, o_2, \dots, o_n with $o_i = a$ for some i ($1 \leq i \leq n$), the b can be inserted at any position j with $i < j \leq n+1$ (for the special case of a condition $a = b$, $i = j$ is also possible). So, an entry of an ordering can only affect entries on the right of its own position. This means that it is unnecessary to consider those parts of an ordering which are behind the length of the longest interesting order; since that part cannot influence any entries relevant for plan generation, it can be omitted. Therefore, the orderings created by functional dependencies can be cut off after the maximum length of interesting orders, which results in less possible combinations and a smaller NFSM.

The space of possible orderings can be limited further by taking into account the prefix of the ordering: before inserting an entry b in an ordering o_1, o_2, \dots, o_n at the position i , check if there is actually an interesting order with the prefix $o_1, o_2, \dots, o_{i-1}, b$ and stop inserting if no interesting order is found. Also limit the new ordering to the length of the longest matching interesting order; further attributes will never be used. If functional dependencies of the form $a = b$ occur, they might influence the prefix of the ordering and the simple test described above is not sufficient. Therefore, a representative is chosen for each equivalence class created by these dependencies, and for the prefix test the attributes are replaced with their representatives. Since the set of interesting orders with a prefix of o_1, \dots, o_n is a superset of the set for the prefix o_1, \dots, o_n, o_{n+1} , this heuristic can be implemented very efficiently by iterating over i and reducing the set as needed.

Additional techniques can be used to avoid creating superfluous artificial states for groupings: First, in Step 2.3 (see Figure 23.6) the set of attributes occurring in interesting groupings is determined:

$$A_G = \{a \mid \exists g \in G_I : a \in g\}$$

Now, for every attribute a occurring on the right-hand side of a functional dependency the set of potentially reachable relevant attributes is determined:

$$\begin{aligned} r(a, 0) &= \{a\} \\ r(a, n) &= r(a, n-1) \cup \\ &\quad \{a' \mid \exists (a_1 \dots a_m \rightarrow a') \in \mathcal{F} : \\ &\quad \quad \{a_1 \dots a_m\} \cap r(a, n-1) \neq \emptyset\} \\ r(a) &= r(a, |\mathcal{F}|) \cap A_G \end{aligned}$$

This can be used to determine if a functional dependency actually adds useful attributes. Given a functional dependency $a_1 \dots a_n \rightarrow a$ and a grouping g with $\{a_1 \dots a_n\} \subseteq g$, a should only be added to g if $r(a) \not\subseteq g$, i.e. the attribute might actually lead to a new interesting grouping. For example, given the interesting groupings $\{a\}$, $\{a, b\}$ and the functional dependencies $a \rightarrow c$, $a \rightarrow d$, $d = b$. When considering the grouping $\{a\}$, the functional dependency $a \rightarrow c$ can be ignored, as it can only produce the attribute c , which does not occur in an interesting grouping. However, the functional dependency $a \rightarrow d$ should be added, since transitively the attribute b can be produced, which does occur in an interesting grouping.

Since there are no ϵ edges between groupings, i.e. groupings are not compatible with each other, a grouping can only be relevant for the query if it is a subset of an interesting ordering (as further attributes could be added by functional dependencies). However, a simple subset test is not sufficient, as equations of the form $a = b$ are also supported; these can effectively rename attributes, resulting in a slightly more complicated test:

In Step 2.3 (see Figure 23.6) the equivalence classes induced by the equations in \mathcal{F} are determined and for each class a representative is chosen (a and $a_1 \dots a_n$ are attributes occurring in the G_I):

$$\begin{aligned} E(a, 0) &= \{a\} \\ E(a, n) &= E(a, n-1) \cup \\ &\quad \{a' \mid ((a = a') \in \mathcal{F}) \vee ((a' = a) \in \mathcal{F})\} \\ E(a) &= E(a, |\mathcal{F}|) \\ e(a) &= \text{a representative chosen from } E(A) \\ e(\{a_1 \dots a_n\}) &= \{e(a_1) \dots e(a_n)\}. \end{aligned}$$

Using these equivalence classes, a mapped set of interesting groupings is produced that will be used to test if a grouping is relevant:

$$G_I^E = \{e(g) \mid g \in G_I\}$$

Now a grouping g can be pruned if $\nexists g' \in G_I^E : e(g) \subseteq g'$. For example, given the interesting grouping $\{a\}$ and the equations $a = b$, $b = c$, the grouping $\{d\}$ can be pruned, as it will never lead to an interesting grouping; however, the groupings $\{b\}$ and $\{c\}$ have to be kept, as they could change to an interesting grouping later on.

Note that although they appear to test similar conditions, the first pruning technique (using $r(a)$) is not dominated by the second one (using $e(a)$). Consider e.g. the interesting grouping $\{a\}$, the equation $a = b$ and the functional dependency $a \rightarrow b$. Using only the second technique, the grouping $\{a, b\}$ would be created, although it is not relevant.

23.4.8 Complex Ordering Requirements

Specifying the ordering requirements of an operator can be surprisingly difficult. Consider the following SQL query:

```

select *
from   S s, R r
where  r.a=s.a and r.b=s.b and
       r.c=s.c and r.d=s.d

```

When answering this query using a sort-merge join, the operator has to request a certain ordering. But there are many orderings that could be used: The intuitive ordering would be $abcd$, but $adcb$ or any other permutation could have been used as well. This is problematic, as checking for an exponential number of possibilities is not acceptable in general. Note that this problem is not specific to our approach, the same is true, e.g., for Simmen's approach.

The problem can be solved by defining a total ordering between the attributes, such that a canonical ordering can be constructed. We give some rules how to derive such an ordering below, but it can happen that such an ordering is unavailable (or rather the construction rules are ambiguous). Given, for example, two indices, one on $abcd$ and one on $adcb$, both orderings would be a reasonable choice. If this happens, the operators have two choices: Either they accept all reasonable orderings (which could still be an exponential number, but most likely only a few orderings remaining) or they limit themselves to one ordering, which could induce unnecessary sort operators. Probably the second choice is preferable, as the ambiguous case should be rare and does not justify the complex logic of the first solution.

The attribute ordering can be derived by using the following heuristical rules:

1. Only attributes that occur in sets without natural ordering (i.e. complex join predicates or grouping attributes) have to be ordered.
2. Orderings that are given (e.g., indices, user-requested orderings etc.) order some attributes.
3. Small orderings should be considered first. If an operator requires an ordering with the attributes abc , and another operator requires an ordering with the attributes bc , the attributes b and c should come before a .
4. The attributes should be ordered according to equivalence classes. If a is ordered before b , all orderings in $E(a)$ should be ordered before all orderings in $E(b)$.
5. Attributes should be ordered according to the functional dependencies, i.e. if $a \rightarrow b$, a should come before b . Note that $a = b$ suggests no ordering between a and b .
6. The remaining unordered attributes can be ordered in an arbitrary way.

The rules must check if they create contradictions. If this happens, the contradicting ordering must be omitted, resulting in potentially superfluous sort operators. Note that in some cases these sort operators are simply unavoidable: If for the example query one index on R exists with the ordering $abcd$ and one

n	#Edges	t (ms)	#Plans	t/plan	t (ms)	#Plans	t/plan	% t	% #Plans	%. t/plan
5	n-1	2	1541	1.29	1	1274	0.78	2.00	1.21	1.65
6	n-1	9	7692	1.17	2	5994	0.33	4.50	1.28	3.55
7	n-1	45	36195	1.24	12	26980	0.44	3.75	1.34	2.82
8	n-1	289	164192	1.76	74	116562	0.63	3.91	1.41	2.79
9	n-1	1741	734092	2.37	390	493594	0.79	4.46	1.49	3.00
10	n-1	11920	3284381	3.62	1984	2071035	0.95	6.01	1.59	3.81
5	n	4	3060	1.30	1	2051	0.48	4.00	1.49	2.71
6	n	21	14733	1.42	4	9213	0.43	5.25	1.60	3.30
7	n	98	64686	1.51	20	39734	0.50	4.90	1.63	3.02
8	n	583	272101	2.14	95	149451	0.63	6.14	1.82	3.40
9	n	4132	1204958	3.42	504	666087	0.75	8.20	1.81	4.56
10	n	26764	4928984	5.42	2024	2465646	0.82	13.22	2.00	6.61
5	n+1	12	5974	2.00	1	3016	0.33	12.00	1.98	6.06
6	n+1	69	26819	2.57	6	12759	0.47	11.50	2.10	5.47
7	n+1	370	119358	3.09	28	54121	0.51	13.21	2.21	6.06
8	n+1	2613	509895	5.12	145	208351	0.69	18.02	2.45	7.42
9	n+1	27765	2097842	13.23	631	827910	0.76	44.00	2.53	17.41
10	n+1	202832	7779662	26.07	3021	3400945	0.88	67.14	2.29	29.62

Figure 23.14: Plan generation for different join graphs, Simmen’s algorithm (left) vs. our algorithm (middle)

index on S with the ordering dcb , the heuristical rules detect a contradiction and choose one of the orderings. This results in a sort operator before the (sort-merge) join, but this sort could not have been avoided anyway.

23.5 Experimental Results

The framework described in this chapter solves two problems: First, it provides an efficient representation for reasoning about orderings and second, it allows keeping track of orderings and groupings at the same time. Since these topics are treated separately in the related work, the experimental results are split in two sections: In Section 23.6 the framework is compared to another published framework while only considering orderings, and in Section 23.7 the influence of groupings is evaluated.

23.6 Total Impact

We now consider how order processing influences the time needed for plan generation. Therefore, we implemented both our algorithm and the algorithm proposed by Simmen et al. [709, 710] and integrated them into a bottom-up plan generator based on [488].

To get a fair comparison, we tuned Simmen’s algorithm as much as possible. The most important measure was to cache results in order to eliminate repeated calls to the very expensive *reduce* operation. Second, since Simmen’s algorithm requires dynamic memory, we implemented a specially tailored memory management. This alone gave us a speed up by a factor of three. We further tuned the algorithm by thoroughly profiling it until no more improvements were possible. For each order optimization framework the plan generator was recompiled to allow for as many compiler optimizations as possible. We also carefully observed that in all cases both order optimization algorithms produced the same optimal plan.

We first measured the plan generation times and memory usage for TPC-R Query 8. A detailed discussion of this query follows in Section 23.7, here

we ignored the grouping properties to compare it with Simmen's algorithm. The result of this experiment is summarized in the following table. Since order optimization is tightly integrated with plan generation, it is impossible to exactly measure the time spent just for order optimization during plan generation. Hence, we decided to measure the impact of order optimization on the total plan generation time. This has the advantage that we can also (for the first time) measure the impact order optimization has on plan generation time. This is important since one could argue that we are optimizing a problem with no significant impact on plan generation time, hence solving a non-problem. As we will see, this is definitely not the case.

In subsequent tables, we denote by $t(ms)$ the total execution time for plan generation measured in milliseconds, by $\#Plans$ the total number of subplans generated, by $t/plan$ the average time (in microseconds) needed to introduce one plan operator, i.e. the time to produce a single subplan, and by $Memory$ the total memory (in KB) consumed by the order optimization algorithms.

	Simmen	Our algorithm
t (ms)	262	52
#Plans	200536	123954
t/plan (μs)	1.31	0.42
Memory (KB)	329	136

From these numbers, it becomes obvious that order optimization has a significant influence on total plan generation time. It may come as a surprise that fewer plans need to be generated by our approach. This is due to the fact that the (reduced) FSM only contains the information relevant to the query, resulting in fewer states. With Simmen's approach, the plan generator can only discard plans if the ordering is the same and the set of functional dependencies is equal (respectively a subset). It does not recognize that the additional information is not relevant for the query.

In order to show the influence of the query on the possible gains of our algorithm, we generated queries with 5-10 relations and a varying number of join predicates—that is, edges in the join graph. We always started from a chain query and then randomly added some edges. For small queries we averaged the results of 100 queries and averaged 10 queries for large queries. The results of the experiment can be found in Fig. 23.14. In the second column, we denote the number of edges in terms of the number of relations (n) given in the first column. The next six columns contain (1) the total time needed for plan generation (in ms), (2) the number of (sub-) plans generated, and (3) the time needed to generate a subplan (in μs), i.e. to add a single plan operator, for (a) Simmen's algorithm (columns 3-5) and our algorithm (columns 6-8). The total plan generation time includes building the DFSM when our algorithm is used. The last three columns contain the improvement factors for these three measures achieved by our algorithm. More specifically, column $\% x$ contains the result of dividing the x column of Simmen's algorithm by the corresponding x column entry of our algorithm.

Note that we are able to keep the plan generation time below one second in most cases and three seconds in the worst case, whereas when Simmen's

n	#Edges	Simmen	Our Algorithm	DFSM
5	n-1	14	10	2
6	n-1	44	28	2
7	n-1	123	77	2
8	n-1	383	241	3
9	n-1	1092	668	3
10	n-1	3307	1972	4
5	n	27	12	2
6	n	68	36	2
7	n	238	98	3
8	n	688	317	3
9	n	1854	855	4
10	n	5294	2266	4
5	n+1	53	15	2
6	n+1	146	49	3
7	n+1	404	118	3
8	n+1	1247	346	4
9	n+1	2641	1051	4
10	n+1	8736	3003	5

Figure 23.15: Memory consumption in KB for Figure 23.14

algorithm is applied, plan generation time can be as high as 200 seconds. This observation leads to two important conclusions:

1. Order optimization has a significant impact on total plan generation time.
2. By using our algorithm, significant performance gains are possible.

For completeness, we also give the memory consumption during plan generation for the two order optimization algorithms (see Fig. 23.15). For our approach, we also give the sizes of the DFSM which are included in the total memory consumption. All memory sizes are in KB. As one can see, our approach consumes about half as much memory as Simmen's algorithm.

23.7 Influence of Groupings

Integrating groupings in the order optimization framework allows the plan generator to easily exploit groupings and, thus, produce better plans. However, order optimization itself might become prohibitively expensive by considering groupings. Therefore, we evaluated the costs of including groupings for different queries.

Since adding support for groupings has no effect on the runtime behavior of the plan generator (all operations are still one table lookup), we measured the runtime and the memory consumption of the preparation step both with and without considering groupings. When considering groupings, we treated each interesting ordering also as an interesting grouping, i.e. we assumed that a grouping-based (e.g. hash-based) operator was always available as an alternative. Since this is the worst-case scenario, it should give an upper bound for

the additional costs. All experiments were performed on a 2.4 GHz Pentium IV, using the gcc 3.3.1.

To examine the impact for real queries, we choose a more complex query from the well-known TPC-R benchmark ([762], Query 8):

```

select
  o_year,
  sum(case when nation = '[NATION]'
        then volume
        else 0
      end) / sum(volume) as mkt_share
from
  (select
    extract(year from o_orderdate) as o_year,
    l_extendedprice * (1-l_discount) as volume,
    n2.n_name as nation
  from part,supplier,lineitem,orders,customer,
       nation n1,nation n2,region
  where
    p_partkey = l_partkey and
    s_suppkey = l_suppkey and
    l_orderkey = o_orderkey and
    o_custkey = c_custkey and
    c_nationkey = n1.n_nationkey and
    n1.n_regionkey = r_regionkey and
    r_name = '[REGION]' and
    s_nationkey = n2.n_nationkey and
    o_orderdate between date '1995-01-01' and
                      date '1996-12-31' and
    p_type = '[TYPE]'
  ) as all_nations
group by o_year
order by o_year;

```

When considering this query, all attributes used in joins, group-by and order-by clauses are added to the set of interesting orders. Since hash-based solutions are possible, they are also added to the set of interesting groupings.

This results in the sets

$$\begin{aligned}
O_I^P &= \{(o_year), (o_partkey), (p_partkey), \\
&\quad (l_partkey), (l_suppkey), (l_orderkey), \\
&\quad (o_orderkey), (o_custkey), (c_custkey), \\
&\quad (c_nationkey), (n1.n_nationkey), \\
&\quad (n2.n_nationkey), (n_regionkey), \\
&\quad (r_regionkey), (s_suppkey), (s_nationkey)\} \\
O_I^T &= \emptyset \\
G_I^P &= \{\{o_year\}, \{o_partkey\}, \{p_partkey\}, \\
&\quad \{l_partkey\}, \{l_suppkey\}, \{l_orderkey\}, \\
&\quad \{o_orderkey\}, \{o_custkey\}, \{c_custkey\}, \\
&\quad \{c_nationkey\}, \{n1.n_nationkey\}, \\
&\quad \{n2.n_nationkey\}, \{n_regionkey\}, \\
&\quad \{r_regionkey\}, \{s_suppkey\}, \{s_nationkey\}\} \\
G_I^T &= \emptyset
\end{aligned}$$

Note that here O_I^T and G_I^T are empty, as we assumed that each ordering and grouping would be produced if beneficial. For example, we might assume that it makes no sense to intentionally group by o_year : If a tuple stream is already grouped by o_year it makes sense to exploit this, however, instead of just grouping by o_year it could make sense to sort by o_year , as this is required anyway (although here it only makes sense if the sort operator performs early aggregation). In this case, $\{o_year\}$ would move from G_I^P to G_I^T , as it would be only tested for, but not produced.

The set of functional dependencies (and equations) contains all join conditions and constant conditions:

$$\begin{aligned}
\mathcal{F} &= \{\{p_partkey = l_partkey\}, \{\emptyset \rightarrow p_type\}, \\
&\quad \{o_custkey = c_custkey\}, \{\emptyset \rightarrow r_name\}, \\
&\quad \{c_nationkey = n1.n_nationkey\}, \\
&\quad \{s_nationkey = n2.n_nationkey\}, \\
&\quad \{l_orderkey = o_orderkey\}, \\
&\quad \{s_suppkey = l_suppkey\}, \\
&\quad \{n1.n_regionkey = r_regionkey\}\}
\end{aligned}$$

To measure the influence of groupings, the preparation step was executed twice: Once with the data as given above and once with $G_I^P = \emptyset$ (i.e. groupings were ignored). The space and time requirements are shown below:

	With Groups	Without Groups
Duration [ms]	0.6ms	0.3ms
DFSM [nodes]	63	32
Memory [KB]	5	2




Figure 23.16: Time requirements for the preparation step

Here time and space requirements both increase by a factor of two. Since all interesting orderings are also treated as interesting groupings, a factor of about two was expected.

While Query 8 is one of the more complex TPC-R queries, it is not overly complex when looking at order optimization. It contains 16 interesting orderings/groupings and 8 functional dependencies, but they cannot be combined in many reasonable ways, resulting in a comparatively small DFSM. In order to get more complex examples, we produced randomized queries with 5 – 10 relations and a varying number of join predicates. We always started from a chain query and then randomly added additional edges to the join graph. The results are shown for $n - 1$, n and $n + 1$ additional edges. In the case of 10 relations, this means that the join graph consisted of 18, 19 and 20 edges, respectively.

The time and space requirements for the preparation step are shown in Figure 23.16 and Figure 23.17, respectively. For each number of relations, the requirements for the combined framework (o+g) and the framework ignoring groupings (o) are shown. The numbers in parentheses ($n - 1$, n and $n + 1$) are the number of additional edges in the join graph.

As with Query 8, the time and space requirements roughly increase by a factor of two when adding groupings. This is a very positive result, given that a factor of two can be estimated as a lower bound (since every interesting ordering is also an interesting grouping here). Furthermore, the absolute time and space requirements are very low (a few ms and a few KB), encouraging the inclusion of groupings in the order optimization framework.




Figure 23.17: Space requirements for the preparation step

23.8 Annotated Bibliography

Very few papers exist on order optimization. While the problem of optimizing interesting orders was already introduced by Selinger et al. [672], later papers usually concentrated on exploiting, pushing down or combining orders, not on the abstract handling of orders during query optimization.

Papers by Simmen, Shekita, and Malkemus [709, 710] introduced a framework based on functional dependencies for reasoning about orderings. Since this is the only paper which really concentrates on the abstract handling orders and our approach is similar in the usage of functional dependencies, we will describe their approach in some more detail.

For a plan node they keep just a single (physical) ordering. Additionally, they associate all the applicable functional dependencies with a plan node. Hence, the lower-bound space requirement for this representation is essentially $\Omega(n)$, where n is the number of functional dependencies derived from the query. Note that the set of functional dependencies is still (typically) much smaller than the set of all logical orderings. In order to compute the function `containsOrdering`, Simmen et al. apply a *reduction algorithm* on both the ordering associated with a plan node and the ordering given as an argument to `containsOrdering`. Their reduction roughly does the opposite of deducing more orderings using functional dependencies. Let us briefly illustrate the reduction by an example. Assume the physical ordering a tuple stream satisfies is (a) , and the required ordering is (a, b, c) . Further assume that there are two functional dependencies available: $a \rightarrow b$ and $a, b \rightarrow c$. The reduction algorithm is performed on both orderings. Since (a) is already minimal, nothing changes. Let us now reduce (a, b, c) . We apply the second functional dependency first. Using $a, b \rightarrow c$, the reduction algorithm yields (a, b) because c appears in (a, b, c)

after a and b . Hence, c is removed. In general, every occurrence of an attribute on the right-hand side of a functional dependency is removed if all attributes of the left-hand side of the functional dependency precede the occurrence. Reduction of (a, b) by $a \rightarrow b$ yields (a) . After both orderings are reduced, the algorithm tests whether the reduced required ordering is a prefix of the reduced physical ordering. Note that if we applied $a \rightarrow b$ first, then (a, b, c) would reduce to (a, c) and no further reduction would be possible. Hence, the rewrite system induced by their reduction process is not confluent. This problem is not mentioned by Simmen et al., but can have the effect that `containsOrdering` returns *false* whereas it should return *true*. The result is that some orderings remain unexploited; this could be avoided by maintaining a minimal set of functional dependencies, but the computation costs would probably be prohibitive. This problem does not occur with our approach. On the complexity side, every functional dependency has to be considered by the reduction algorithm at least once. Hence, the lower time bound is $\Omega(n)$.

In case all functional dependencies are introduced by a single plan node and all of them have to be inserted into the set of functional dependencies associated with that plan node, the lower bound for `inferNewLogicalOrderings` is also $\Omega(n)$.

Overall, Simmen et al. proposed the important framework for order optimization utilizing functional dependencies and nice algorithms to handle orderings during plan generation, but the space and time requirements are unfortunate since plan generation might generate millions of subplans. Also note that the reduction algorithm is not applicable for groupings (which, of course, was never intended by Simmen): Given the grouping $\{a, b, c\}$ and the functional dependencies $a \rightarrow b$ and $b \rightarrow c$, the grouping would be reduced to $\{a, c\}$ or to $\{a\}$, depending on the order in which the reductions are performed. This problem does not occur with orderings, as the attributes are sorted and can be reduced back to front.

A recent paper by Wang and Cherniack [790] presented the idea of combining order optimization with the optimization of groupings. Based upon Simmen's framework, they annotated each attribute in an ordering with the information whether it is actually ordered by or grouped by. For a single attribute a , they write $O_{aO}(R)$ to denote that R is ordered by a , $O_{aG}(R)$ to denote that R is grouped by a and $O_{aO \rightarrow bG}$ to denote that R is first ordered by a and then grouped by b (within blocks of the same a value). Before checking if a required ordering or grouping is satisfied by a given plan, they use some inference rules to get all orderings and groupings satisfied by the plan. Basically, this is Simmen's reduction algorithm with two extra transformations for groupings. In their paper the check itself is just written as \in , however, at least one reduction on the required ordering would be needed for this to work (and even that would not be trivial, as the stated transformations on groupings are ambiguous). The promised details in the cited technical report are currently not available, as the report has not appeared yet. Also note that, as explained above, the reduction approach is fundamentally not suited for groupings. In Wang's and Cherniack's paper, this problem does not occur, as they only look at a very specialized kind of grouping: As stated in their Axiom 3.6, they assume that a grouping

$O_{aG \rightarrow bG}$ is first grouped by a and then (within the block of tuples with the same a value) grouped by b . However, this is a very strong condition that is usually not satisfied by a hash-based grouping operator. Therefore, their work is not general enough to capture the full functionality offered by a state-of-the-art query execution engine.

In this chapter, we followed [544, 543].

Chapter 24

Cardinality and Cost Estimates

24.1 Introduction

The plan generator relies on a cost function to evaluate the different plans and to determine the cheapest one. This chapter is concerned with the development of cost functions. The main input to cost functions are cardinalities. For example, assume a scan of a relation, which also applies a selection predicate. Clearly, the cost of scanning the relation depends on the physical layout of the relation on disk. Further, the CPU cost for evaluating the predicate depends on the number of tuples in the relation. Note that the cardinality of a relation is independent of its physical layout.

In general, the cost of an algebraic operator is estimated by using a *profile* of the database. The profile must be small, e.g. a couple of kilobytes per relation¹. According to the above lines, we distinguish between the *logical* and *physical* profile. For each database item and its constituents, there exist specialized logical and physical profiles. They exist for relations, indices, attributes, and sets of attributes. Fig. ?? shows the typical classes and their associations for the different profiles. The minimal logical profile of a relation R is its cardinality $|R|$. Its minimal physical profile consists of the number of pages $\lceil |R|/p \rceil$ on which it is stored. In Chapter 4, we saw more advanced physical profiles. ToDo

The DBMS must be capable to perform several operations to derive profiles and to deal with them. Fig. 24.1 gives an overview. This figure roughly follows the approach of Mannino et al. [501, ?]. The first operation is the *build* operation, that takes as input a specification of the profiles to be build (because there are many different alternatives, as we will see) and the database. From that, it builds the according profiles for all database items of all the different granularities. When updates arrive, the profiles must be updated. This can either be done by a complete recalculation or by an incremental update operation on the profiles themselves. The latter is reflected in the operation **update**. Unfortunately, not all profiles can have an update operation. Within this book,

¹Given today's cost for main memory, it may also be reasonable to use a couple of megabytes.

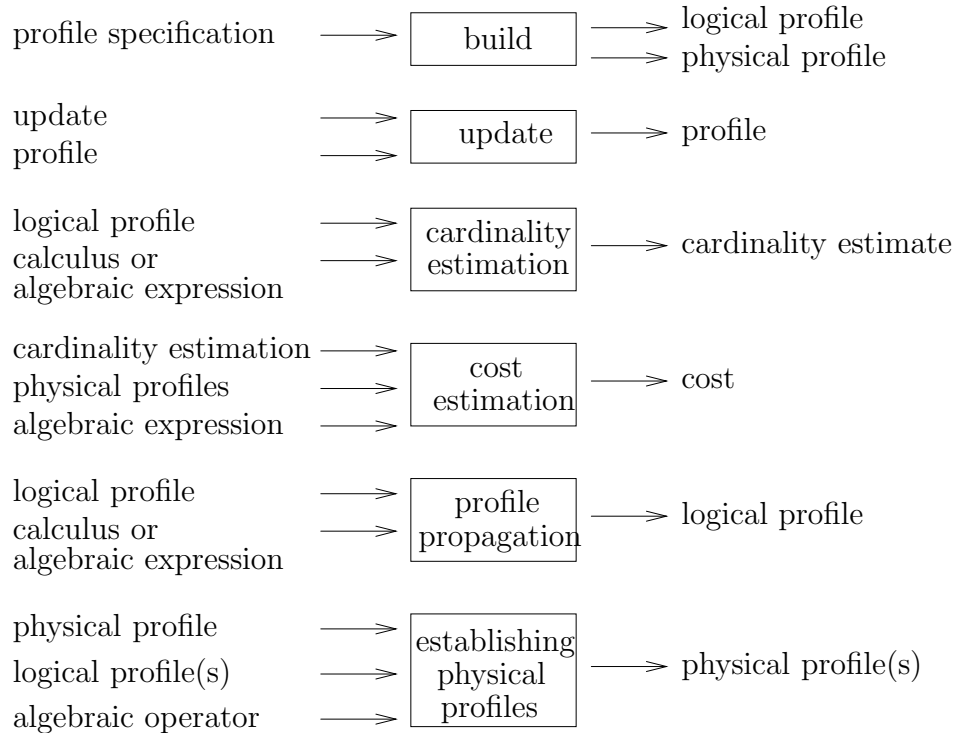


Figure 24.1: Overview of operations for cardinality and cost estimations

we will not be too concerned about building and updating profiles. At the end of this chapter, we will provide some references.

The main operations this chapter deals with are among the remaining ones. The first of them is *cardinality estimation*. Given an algebraic expression or a calculus expression together with a logical profile of the database, we estimate the output/result cardinality of the expression. Why do we say algebraic or calculus expression? Remember that plan generators generate plans for plan classes. Each plan class corresponds to a set of equivalent plans. They all produce the same result and, hence, the same number of output tuples. Thus, one arbitrary representative of the class of equivalent algebraic expressions should suffice to calculate the logical profile, as a logical profile depends only on the outcome. On the other hand, the plan class more directly corresponds to a calculus expression. Hence, estimating the result cardinality of a calculus expression is a viable alternative. In the literature, most papers deal with the first approach while only a few deal with the latter (e.g. [194]).

The second operation we will be concerned with is *cost estimation*. Given logical and physical profiles for all inputs and an algebraic operator (tree), this operation calculates the actual costs.

The third major task is profile propagation. Given a logical or physical profile and an expression, we must be able to calculate the profile of the result, since this may be the input to other expressions and thus be needed for further cardinality estimates. The estimation of a physical profile occurs mostly in cases where operators write to disk. Given Chapter 4, this task is easy enough

to be left to the reader.

Since we follow the algebraic approach, we must be able to calculate the output cardinality of every operator occurring in the algebra. This task is vastly simplified by observing the following, where we have been quite sloppy with the subscripts of the operations.

$$\begin{aligned}
|\chi_{a:e_2}(e_1)| &= |e_1| \\
|\Gamma_{g,f}(e)| &= |\Pi_g^D(e)| \\
|e_1 \Gamma e_2| &= |e_1| \\
|e_1 \triangleright e_2| &= |e_1| - |e_1 \bowtie e_2| \\
|e_1 \bowtie e_2| &= |e_1 \bowtie e_2| + |e_1 \triangleright e_2| \\
|e_1 \bowtie e_2| &= |e_1 \bowtie e_2| + |e_1 \triangleright e_2| + |e_2 \triangleright e_1| \\
|e_1 \bowtie e_2| &= |\Pi_{\mathcal{A}(e_1) \cup \mathcal{A}(e_2)}^D(e_1 \bowtie e_2)| \\
|\text{Sort}(e)| &= |e| \\
|\text{Tmp}(e)| &= |e| \\
|e_1 \times e_2| &= |e_1| * |e_2| \\
|\Pi_A(e)| &= |e| \quad (\text{bag semantics}) \\
|e_1 \cup e_2| &= |\Pi_{\mathcal{A}(e_1)}^D(e_1 \bar{\cup} e_2)| \quad (\text{bag vs. set semantics}) \\
|e_1 \cap e_2| &= |e_1 \bowtie e_2| \quad \text{equijoin over all attributes} \\
|e_1 \setminus e_2| &= |e_1| - |e_1 \cap e_2| \\
|e_1 \bar{\cup} e_2| &= |e_1| + |e_2| \quad \text{bag semantics} \\
|\Pi_{\alpha \cup \beta}^D(R)| &= |\Pi_{\alpha}^D(R)| \quad \text{if there is an FD } \alpha \rightarrow \beta
\end{aligned}$$

This shows that we can go a far way if we are able to estimate the output cardinality for duplicate eliminating projections, selections, (bag) joins, and semijoins. For a certain class of profiles, Richard shows that a profile consisting ‘only’ of the sizes of all duplicate eliminating projections on all subsets of attributes of all relations is a complete profile under certain assumptions [621]. Since the set of subsets of a set of attributes can be quite large, Richard exploits functional dependencies to reduce this set by exploiting the fact that $|\Pi_{\alpha \cup \beta}^D(R)| = |\Pi_{\alpha}^D(R)|$ if there exists a functional dependency $\alpha \rightarrow \beta$.

A major differentiator for logical attribute profiles is the kind of the domain of the attribute. We distinguish between categorical attributes (e.g. color), discrete ordered domains (e.g. integer attributes, decimals, strings), and continuous ordered domains (e.g. float). Categorical domains may be ordered or unordered. In the first case they are called *ordinal*, in the latter *nominal*. We will be mainly concerned with integer attributes. Strings are special, and we discuss some approaches in Sec. ???. Continuous domains are also special. The probability of occurrence of any value in a continuous domain in a finite set is zero. The techniques developed in this section can often easily be adopted to continuous domains, even if we do not mention this explicitly.

24.2 A First Approach

The first approach to cost and cardinality estimation integrated into a dynamic programming based plan generator was presented by Selinger et al. [672]. We will use it as the basis for this section.

24.2.1 Top-Most Cost Formula (Overall Costs)

Their top-most cost formula states that the total cost of a query evaluation plan equals the weighted sum of the I/O and CPU costs:

$$\mathcal{C} = \mathcal{C}_{I/O} + w\mathcal{C}_{cpu} \quad (24.1)$$

where w is the weight which can be adapted to different situations. If, for example, the system is CPU bound, we should increase w and if it is I/O bound, we decrease w .

However, it is not totally clear what we are going to optimize under this cost formula. One interpretation could be the following. Assume $w = 0.5$. Then, we could interpret the total costs as response time under the assumption that fifty percent of the CPU time can be executed in parallel with I/O. Accordingly, we find other top-most cost formulas. For example, the weight is sometimes dropped [351]:

$$\mathcal{C} = \mathcal{C}_{I/O} + \mathcal{C}_{cpu} \quad (24.2)$$

Under the above interpretation, this would mean that concurrency is totally absent. The opposite, total concurrency between I/O and CPU, can also be found [144]:

$$\mathcal{C} = \max(\mathcal{C}_{I/O}, \mathcal{C}_{cpu}) \quad (24.3)$$

In these green days, an alternative is to calculate the power consumption during query execution.

24.2.2 Summation of Operator Costs

Given a query evaluation plan, the task is to calculate its I/O and CPU costs. This can be done by calculating the costs for each operator (op) occurring in the query evaluation plan (QEP) and adding up the according costs:

$$\begin{aligned} \mathcal{C}_{I/O} &= \sum_{op \in QEP} \mathcal{C}_{I/O}(op) \\ \mathcal{C}_{cpu} &= \sum_{op \in QEP} \mathcal{C}_{cpu}(op) \end{aligned}$$

However, these formula sometimes raise a problem. For example, the nested loop join method requires multiple evaluations of its inner part.

Further, in order to count the I/O cost correctly, it is necessary to make some assumptions when intermediate results are written to and read from disk. We will use the following assumption: Every operator is responsible for passing its result to the next operator via main memory. For example, a sort merge

join may require sorting its inputs. The sort operators are then responsible for handing over their result to the merge join via main memory. This means that the merge join may not require any I/O if the merge can be done purely in main memory.

24.2.3 CPU Cost

To get an estimate of the CPU costs, Selinger et al. simply count the number of calls to the tuple oriented interface (called RSI). This roughly corresponds to the number of `next` calls in an iterator-based implementation of algebraic operators (for details on the Research Storage Engine (RSS) and its interface RSI see [37]). Hence, what needs to be known is the output cardinality of each operator in the query plan.

Given the input and output cardinalities, it is often quite straightforward to calculate the CPU costs. Take, for example, a selection operator. Clearly, the selection predicate is called n times if n is the input cardinality. The selection predicate itself consists of several calls to comparison functions, Boolean operators, arithmetic operators and the like. The CPU cost of each of these operators can easily be determined (by counting CPU cycles or measurements), and thus the total CPU cost of a selection operator can be determined easily. Other operators are also straightforward. A problem only arises if functions are called whose CPU costs can not easily be determined since they depend on their parameters. A typical example are string comparisons, where the CPU costs depend on the length of the string. Another example are user-defined functions. The framework presented in [?] can be used for all more complex functions. Another possibility is to use simplifying assumptions. The functions we are talking about are executed on a per tuple basis. As there are typically many tuples, using the average execution time for cost calculations is not a bad idea.

24.2.4 Abbreviations

We need some abbreviations to state our cost formulas. A first bunch of them is summarized in Table 24.1. There, we assume that an index is always a B^+ tree.

24.2.5 I/O Costs

Selinger et al. measure I/O costs in the *number of pages read*. Let us first discuss the different possible access paths to a single relation. Clearly, for a simple scan of a relation R , $\|R\|$ pages have to be read. The next access path is composed of an access to a non-clustered index I to retrieve the tuple identifiers of those tuples that satisfy a predicate p followed by an access to the base relation R . Let $F(p)$ be the fraction of tuples satisfying a certain predicate p . $F(p)$ is called the *selectivity* of p . It is the main focus of the next subsection. Selinger et al. distinguish two cases. In the first case, all pages containing qualifying tuples fit into main memory. For this case, they estimate the number of pages accessed

R,S,T	relations
I	index
A,B,C	attributes or sets of attributes
D_A	$\Pi_A^D(R)$
d_A	$ D_A $
\min_A	$\min \Pi_A^D(R)$ for an attribute A of R
\max_A	$\max \Pi_A^D(R)$ for an attribute A of R
$ R $	number of tuples of R
$\ R\ $	number of pages on which R is stored
$\ A\ _B$	average length of a value of attribute A of R (in bytes)
$\ \mathcal{A}(R)\ _B$	average length of a tuple in bytes
$\ I\ $	number of leaf pages of an index
$H(I)$	depth of the index I minus 1

Table 24.1: Notational conventions

by

$$H(T) + F(p) * (\|I\| + \|R\|).$$

EXC

Note that with the help of Chapter 4, we can already do better. In the second case, where the pages containing qualifying tuples do not fit into main memory, they give the estimate

$$H(T) + F(p) * (\|I\| + |R|)$$

for the number of pages read.

In case of a clustered index, they estimate the number of pages read by

$$H(T) + F(p) * (\|I\| + \|R\|).$$

Next, we have to discuss the costs of different join methods. Selinger et al. propose cost formulas for the simple nested loop join (\bowtie^{nl}) and the sort merge join (\bowtie^{sm}). Since summing up the costs of all operators in a tree results in some problems for nested loop joins, they adhere to a recursive computation of total costs. Let e_1 and e_2 be two algebraic expressions. Then they estimate the cost of the simple nested loop join as follows:

$$C_{I/O}(e_1 \bowtie^{nl} e_2) = C_{I/O}(e_1) + |e_1| * C_{I/O}(e_2)$$

where $|e_1|$ denotes the number of tuples produced by the expression e_1 .

As the cost calculation for the sort merge join is not convincing, we follow our own very simple approach here. We split the costs into the sort costs and the merge costs. Given today's memory sizes, it is not unlikely that we need a single merge phase. Hence, the I/O cost for sorting consists of writing and reading the result of e_i if it needs to be sorted. This can be estimated as

$$C_{I/O}(\text{sort}(e_i)) = C_{I/O}(e_i) + 2 * \lceil 1.2 \| \mathcal{A}(e_i) \|_B * |e_i| / \text{pagesize} \rceil$$

where **pagesize** is the page size in bytes. The factor 1.2 is called the *universal fudge factor*. In the above case, it takes care of the storage overhead incurred

by using slotted pages. If we assume that the merge phase of the sort merge join can be performed in main memory, no additional I/O costs occur and we are done.

Clearly, in the light of Chapter 4, counting the numbers of pages read is not sufficient as the discrepancy between random and sequential I/O is tremendous. Thus, better cost functions should use a more elaborate I/O cost model along the lines of Chapter 4. In any case, note that the calculation of the I/O or CPU costs of any operator highly depends on its input and output cardinalities.

24.2.6 Cardinality Estimates

Given a predicate p , we want to estimate its selectivity, which is defined as the fraction of qualifying tuples. If p is a selection predicate applied to a relation R , the selectivity of p is defined as

$$s(p) = \frac{|\sigma_p(R)|}{|R|}.$$

If we know the selectivity of p , we can then easily calculate the result size of a selection:

$$|\sigma_p(R)| = s(p)|R|$$

Similarly for joins. Given a join predicate p and two relations R and S , we define the selectivity of p as

$$s(p) = \frac{|R \bowtie_p S|}{|R \times S|} = \frac{|R \bowtie_p S|}{|R| * |S|}$$

and can calculate the result size of a join by

$$|R \bowtie_p S| = s(p) |R| |S|.$$

The idea of the approach of Selinger et al. is to calculate the result cardinality for a plan class by the following procedure. First, the sizes of all relations represented by the plan class are multiplied. This is the result of their cross product. In a second step, they take a look at the predicate p applied to the relations in the plan class. For p they calculate a selectivity estimate $s(p)$ and multiply it with the result of the first step. This then gives the result. Hence, if a plan class represents the algebraic expression

$$\sigma_p(\times_{i=1}^n R_i),$$

the cardinality estimate is

$$s(p) \prod_{i=1}^n |R_i|.$$

Since p can be a complex predicate involving boolean operators, they have to deal with them. Table 24.2 summarizes the proposed selectivity estimation. A and B denote attributes, c , c_1 , c_2 denote constants, L denotes a list of values, Q denotes a subquery. In System R, the number of distinct values for an attribute

predicate	$s(p)$	comment
$\text{not}(p_1)$	$1 - s(p_1)$	
$p_1 \wedge p_2$	$s(p_1) * s(p_2)$	independence
$p_1 \vee p_2$	$s(p_1) + s(p_2) - s(p_1)s(p_2)$	
$A = c$	$1/d_A$	if d_A is known, uniformity
	$1/10$	else
$A = B$	$1/\max(d_A, d_B)$	if d_A and d_B are known, uniformity
	$1/d_X$	if only $d_X, X \in \{A, B\}$ is known
	$1/10$	else
$A > c$	$\frac{\max_A - c}{\max_A - \min_A}$	if min and max are known, uniformity
	$1/3$	else
$c_1 \leq A \leq c_2$	$\frac{c_2 - c_1}{\max_A - \min_A}$	if min and max are known, uniformity
	$1/4$	else
$A \text{ IN } L$	$\min(1/2, s(A = c) L)$	
$A \text{ IN } Q$	$ Q * X $	X is cross product of all relations in Q 's from clause

Table 24.2: Selectivity estimation as proposed by Selinger et al.[672]

(d_A, d_B) is only known if there exists an according index on the attribute. Let us give some rationale for the selectivity estimation of $A \text{ IN } Q$ for an attribute A and a subquery Q . Assume that A is an attribute of relation R and the subquery Q is of the form `select B from S ...`. Further assume that $\Pi_A(R) \subseteq \Pi_B(S)$, i.e. referential integrity holds. Clearly, if all tuples of S are in the result of Q , the selectivity is equal to 1. If the output cardinality of Q is restricted by a factor $s' = |Q|/|S|$, then we may assume that the number of distinct values in Q 's result is restricted by the same factor. Hence, the selectivity factor of the total predicate is also s' . Selinger et al. now continue as follows: "With a little optimism, we can extend this reasoning to include subqueries which are joins and subqueries in which column [B] is replaced by an arithmetic expression involving column names. This leads to the formula given above."

Discussion Taking a broad view at the above model, we see that

- the estimates for CPU and I/O times are quite rough,
- the approach is not complete, especially projection and semijoin are not treated,
- profile propagation is not discussed.

Further, the uniformity and independence assumptions are applied. This has been shown to be quite inaccurate in many cases. More specifically, applying these and other assumptions often leads to an overestimate of real result cardinalities [151, ?].

How bad is it in terms of plan generation if we under- or overestimate the cardinalities of intermediate results? As Ioannidis and Christodoulakis pointed

out, errors propagate multiplicatively through joins [384]. Assume we want to join eight relations R_1, \dots, R_8 and that the cardinality estimates of R_i are each a factor of 5 off. Then the cardinality estimation of $R_1 \bowtie R_2 \bowtie R_3$ will be a factor of 125 off. Clearly, this can affect the subsequent join ordering. If we were only a factor of 2 off, the cardinality estimation of $R_1 \bowtie R_2 \bowtie R_3$ could be only a factor of eight off. This shows that minimizing the multiplicative error is a serious intention.

The effect of misestimating cardinalities on plan quality has not been thoroughly investigated. There exists a study by Kumar and Stonebraker, which concludes that it does not matter [444]. However, we do not trust this conclusion. In the distributed context, it has been shown by XXX that wrong cardinality estimates lead to bad plans [?]. This may also be true in the central case. Swami and Schiefer give a query and its profiles for which bad cardinality estimates lead to a very bad plan [751]. Later, we will give two further examples showing that good cardinality estimation is vital for generation of good plans (see Sec. ??). Hence, we are very sure that accurate estimation is vital for plan generation. We suggest to the reader to find examples, using the simple C_{out} cost function, where wrong cardinality estimates lead to bad plans. ToDo
EXC

24.3 A First Logical Profile and its Propagation

We call a logical profile *complete* if it allows us to perform cardinality estimation and logical profile propagation for all algebraic operators. In this section, we present an almost complete logical profile and describe the procedure of profile propagation. The main components are the cumulated frequency, i.e. the number of tuples, and the number of distinct values for each attribute in a relation. It is easy to see that we cannot do without either of them. Further, an upper and lower bound for values of an attribute is needed. Again, we will see that we cannot do without them. Hence, the following profile is minimal.

24.3.1 The Logical Profile

For every attribute A of a relation, we define its *logical profile* as a four tuple

$$b_A = [l_A, u_A, f_A, d_A]$$

where l_A is a lower and u_A is an upper bound for the values of A . Further, f_A is the *cumulated frequency*, i.e. the number of tuples with an A value within the bounds, and d_A is the *number of distinct values* occurring as A 's values within the given bounds.

For the purpose of this section, we can define

$$\begin{aligned} l_A &= \min(\Pi_A(R)) \\ u_A &= \max(\Pi_A(R)) \\ f_A &= |R| \\ d_A &= |\Pi_A^D(R)| \end{aligned}$$

If the attribute A is implicit from the context or does not matter, we may omit it in the subscript.

24.3.2 Assumptions

The first two assumptions we make are:

1. All attribute values are uniformly distributed, and
2. the values of all attributes are drawn independently.

Other assumptions will follow.

Often in the formulas developed below, we talk about the domain of the attributes. This is the potential set of values from which a given attribute takes its values. In case of integer attributes, it is easy to see that the domain of attribute A is $[l_A, u_A]$. The size of the domain, denoted by n_A , then is $n_A = u_A - l_A + 1$. For real values, the size is (almost) infinite. Thus, only some of the formulas given below may carry over to attributes whose type is real.

Let us take a closer look at the assumptions. The uniform distribution assumption means that every value occurs about the same number of times. However, this cannot mean that every value of the domain does so, since d_A may be much smaller than n_A . Hence, we refine the *uniform distribution assumption* (UDA) by assuming that every distinct value occurs about f_A/d_A times,

We still need another assumption: the *equal spread assumption* [595]. It is used to answer the question where the occurring distinct values are in the domain. The equal spread assumption states that they occur at equal distance. Let us elaborate a little on this.

For integers, we know the number n_A of possible values from which A can be drawn. It is $n_A = u_A - l_A + 1$. Let us assume that we have only a few distinct values, that is $d_A \ll n_A$. This is not strictly necessary but is good for our intuition. We can now define the spread between two values occurring in A . Let $D_A = \Pi_A^D(R) = \{x_1, \dots, x_{d_A}\}$ where $x_i < x_{i+1}$ be the sorted set of values occurring for attribute A , also known as *active domain*. Then we can define the spread as

$$\Delta_i = x_{i+1} - x_i$$

The *equal spread assumption* (ESA) states that $\Delta_i = \Delta_j$ for all $1 \leq i, j \leq d_A$. Denote this value by Δ_A .

There are three subtypes of the equal spread assumption, depending on whether we assume the lower and upper bounds l_A and u_A belong to D_A . Type I assumes $l_A, u_A \in D_A$. Then Δ_A becomes $(u_A - l_A)/(d_A - 1)$. In case of type II, where $l_A \in D_A$ and $u_A \notin D_A$ holds, we have $\Delta_A = (u_A - l_A)/d_A$. For type III, where $l_A \notin D_A$ and $u_A \in D_A$ we get $\Delta_A = (u_A - l_A)/(d_A + 1)$. As an example, take $l_A = 1$, $u_A = 13$, and $d_A = 3$. Then for the three types we have the different values $12/2 = 6$, $12/3 = 4$, and $12/4 = 3$. It should be clear that the difference is small if d_A is sufficiently large. If d_A is small, we can store the frequency of each value explicitly. Otherwise, it is large, and it does not matter which type we use. We will mainly use type II.

Given the above assumptions (and one more to come), the task is to establish the operations *cardinality estimation* and *logical profile propagation*. The latter implies that we can calculate the logical profile of all attributes of any result relation established by applying some algebraic operator. Assume we have solved this task. Then it is clear that the cumulated frequency f_A , which equals $|R|$ in this section, solves the task of cardinality estimation. Hence, we consider this as done. The use of the cumulated frequency f_A instead of the seemingly simpler cardinality notation $|R|$ is motivated by the fact that a single attribute will have multiple profiles if histograms are applied. To make the formulas of this section readily available for histogram use is the main motivation for using the cumulated frequency.

24.3.3 Propagation for Selections

We start with the selection operation. Let R be a relation and $A, C \in \mathcal{A}(R)$ be two attributes of R . We are given the profiles $b_A = [l_A, u_A, f_A, d_A]$ and $b_C = [l_C, u_C, f_C, d_C]$ and have to calculate the profiles $b'_A = [l'_A, u'_A, f'_A, d'_A]$ and $b'_C = [l'_C, u'_C, f'_C, d'_C]$ for $\sigma_{p(A)}(R)$ for various selection predicates $p(A)$ in attribute A . We assume that the attribute values of A and C are uniformly distributed and that A and C are independent. If a selection predicate uses two attributes A and B , we again need to give the profile propagation for all attributes C , which are different from them.

Exact match queries The first case we consider is $\sigma_{A=c}$ for a constant c . Clearly, $l'_A = c, u'_A = c$. Further,

$$d'_A = \begin{cases} 1 & \text{if } c \in \Pi_A(R) \\ 0 & \text{else} \end{cases}$$

We cannot be sure whether the first or second case occurs. Since no reasonable cardinality estimation should ever return zero, we always assume $c \in \Pi_A(R)$. More generally, we assume that all constants in a query are contained in the database in the according attributes.

As every distinct value occurs about f_A/d_A times, we conclude that $f'_A = f_A/d_A$. A special case occurs if A is the key. Then, we can immediately conclude that $f'_A = 1$.

Let us now consider another attribute $C \in \mathcal{A}(R), C \neq A$. Since $f'_C = f'_A$, we only need to establish d'_C . For the lack of any further knowledge, we keep the lower and upper bounds, i.e. $l'_C = l_C$ and $u'_C = u_C$. To derive the number of distinct values remaining for attribute B , we can use the formula by Yao/Waters (see Sec. 4.16.1) Denote by $s(p) = |\sigma_{A=c}(R)|/|R| = f'_A/f_A$ the fraction of tuples that survives the selection with predicate $p \equiv A = c$. Fix a distinct value for C . Using the uniform distribution assumption, it occurs in f_C/d_C tuples of R . Then, for this value we have $\binom{f_A - f_C/d_C}{f'_A}$ possibilities to chose f'_A tuples without it. The total number of possibilities to chose f'_A tuples is $\binom{f_A}{f'_A}$. Thus, we may conclude that

$$d'_C = d_C \mathcal{Y}_{f'_C/d_C}^{f'_A}(f'_A)$$

Alternatively, we could use

$$d'_C = d_C * (1 - (1 - s(p))^{f_C/d_C})$$

or any other good approximation (see Section 4.16.1).

Range queries Let us now turn to range queries, i.e. selection predicates of the form $c_1 \leq A \leq c_2$, where $l_A \leq c_1 < c_2 \leq u_A$. In all of them, the lower and upper bounds are given by the range, i.e. $l'_A = c_1$ and $u'_A = c_2$. Using the System R approach, we can estimate

$$\begin{aligned} f'_A &= \frac{c_2 - c_1}{u_A - l_A} * f_A \\ d'_A &= \frac{c_2 - c_1}{u_A - l_A} * d_A \end{aligned}$$

This estimate is good for real values.

We could also rewrite the above estimate for the number of distinct values d'_A to

$$d'_A = \frac{c_2 - c_1}{\Delta_A}$$

As soon as we have estimated the number of distinct values in a given range, we can easily derive the cumulated frequency, as every distinct value occurs as often as it did in R . Thus $f'_A = f_A * (d'_A/d_A)$.

For another attribute C , $C \neq A$, the profile propagation is the same as in the case for $A = c$. We only need to define $s(p) = |\sigma_{c_1 \leq A \leq c_2}(R)|/|R|$.

Equality-based correlation The next case we consider is a predicate of the form $A = B$. If $u_A < l_B$ or $u_B < l_A$, the result is empty. If $l_A \neq l_B$ or $u_A \neq u_B$, we first apply a selection with predicate $\max(l_A, l_B) \leq A \leq \min(u_A, u_B)$ and $\max(l_A, l_B) \leq B \leq \min(u_A, u_B)$. So assume w.l.o.g. that $l_A = l_B$ and $u_A = u_B$. Note that $f_A = f_B$. Denote this number by f . Define n to be the number of values in the domain of attributes A and B . For integers, this number is $n = u_A - l_A + 1$. To refer to the elements of the domain, we assume that it is $\{x_1, \dots, x_n\}$ with $x_i < x_{i+1}$.

Let x be a value in the domain. Then we say that R has a *hole* at x in attribute A , if $x \notin \Pi_A(R)$. Consider a value x in the domain. The probability of not having a hole at x in A is

$$p(x \in A) = \frac{\binom{n-1}{d_A-1}}{\binom{n}{d_A}} = \frac{d_A}{n}$$

In general, we have

$$f'_A = f'_B = \sum_{i=1}^n \overline{f_{AP}}(x_i = A) p(x_i = B | x_i = A) \quad (24.4)$$

where $\overline{f_A} = f/d_A$ is the average frequency of a distinct value in $\Pi_A(R)$, $p(x_i = A) = d_A/n$ is the probability that a tuple has x_i as its value for attribute A , and $p(x_i = B|x_i = A)$ is the conditional probability that a tuple has x_i in its B value if it is known that it has an A value x_i .

Let us first consider the special case where $\Pi_A(R) \subseteq \Pi_B(R)$. Then $p(x_i = B|x_i = A)$ becomes $1/d_B$. Hence,

$$f'_A = f'_B = \sum_{i=1}^n \frac{f}{d_A} \frac{d_A}{n} \frac{1}{d_B} = f/d_B$$

For $\Pi_B(R) \subseteq \Pi_A(R)$, we get $f'_A = f'_B = f/d_A$. Summarizing these cases, we may conclude that

$$f'_A = f'_B = \frac{f}{\max(d_A, d_B)}$$

which is the formula applied in System R if indices exist on A and B . Clearly, we can calculate an upper bound on the number of distinct values as

$$d'_A = d'_B = \min(d_A, d_B).$$

Let us estimate the cumulated frequency after the selection if none of the above conditions hold and independence of A and B holds. Then, the conditional probability $p(x_i = B|x_i = A)$ becomes $p(x_i = B) = 1/n$. Thus

$$f'_A = f'_B = \sum_{i=1}^n \frac{f}{d_A} \frac{d_A}{n} \frac{1}{n} = \frac{f}{n}$$

If A and B are independent and uniformly distributed, the number of distinct values $d'_A = d'_B$ can be estimated as follows. According to Section 4.16.1, we can estimate the number of distinct values in $\Pi_{AB}(R)$ as $\mathcal{D}(n*n, |R|)$, where $|R| = f_A = f_B$. Since out of the $n*n$ possible pairs of values only n are of the form (x_i, x_i) , only $n/(n*n) = 1/n$ tuples are of the qualifying form. Using this factor, we derive

$$d'_A = d'_B = \frac{\mathcal{D}(n*n, f_A)}{n}$$

In case of $\Pi_A(R) \subseteq \Pi_B(R)$, only d_A such pairs out of d_A*d_B exist. Thus, the factor becomes $d_A/(d_A*d_B) = 1/d_B$. For $\Pi_B(R) \subseteq \Pi_A(R)$, we have the factor $1/d_A$. Both cases can be summarized as in

$$d'_A = d'_B = \frac{\mathcal{D}(n*n, f_A)}{\max(d_A, d_B)}$$

In case the domain size n is not available, we could estimate it by $|\Pi_A^D(R) \cup \Pi_B^D(S)|$. If this number is not available either, we could hesitatingly use $d_A d_B$.

An alternative is to use

$$d'_A = d'_B = d_A * \mathcal{Y}_{f_A/d_A}^{f_A}(f'_A)$$

or some of its approximations like

$$d'_A = d'_B = d_A * (1 - (1 - s(A = B))^{f_A/d_A}),$$

where $s(A = B) = f'_A/f_A$.

Inequality-based correlation As a last exercise, let us calculate the profile for selections of the form $\sigma_{A \leq B}(R)$. For simplicity, we assume $l_A = l_B$ and $u_A = u_B$. Thus, $l'_A = l'_B = l_A$ and $u'_A = u'_B = u_A$. To calculate the cumulative frequency of the result under independence of A and B , we apply the type I equal spread assumption, with $\Delta_A = (u_A - l_A)/(d_A - 1)$. Hence, we assume that $x_i = l_A + (i - 1)\Delta_A$. This gives us

$$\begin{aligned}
 f'_A &= \sum_{i=1}^{d_A} \overline{f_{AP}}(x_i \leq B | x_i = A) \\
 &= \sum_{i=1}^{d_A} \overline{f_{AP}}(x_i \leq B) \\
 &= \overline{f_A} \sum_{i=1}^{d_A} \frac{x_i - l_B}{u_B - l_B} \\
 &= \overline{f_A} \frac{1}{u_B - l_B} \left(\sum_{i=1}^{d_A} x_i - d_A l_B \right) \\
 &= \overline{f_A} \frac{d_A}{u_B - l_B} \left(l_A - l_B + \Delta_A \frac{d_A - 1}{2} \right) \\
 &= \overline{f_A} \frac{d_A}{u_B - l_B} \frac{(u_A - l_A)(d_A - 1)}{(d_A - 1) \cdot 2} \\
 &= \overline{f_A} \frac{u_A - l_A}{u_B - l_B} \frac{d_A}{2} \\
 &= \frac{\overline{f_A} d_A}{d_A \cdot 2} \\
 &= \frac{\overline{f_A}}{2}
 \end{aligned}$$

As an exercise the reader may verify that $f'_A = (d_A - 1)f_A/(2d_A)$ under the type II equal spread assumption. As an additional exercise the reader should derive d'_A and d'_B . We conjecture that

EXC

$$d'_A = \mathcal{D}(n_A, f'_A)$$

or

$$d'_A = d_A * \mathcal{Y}_{f_A/d_A}^{f'_A}(f'_A).$$

The following observation is crucial: Even if in the original relations the values of A and B are uniformly distributed, which typically is not the case, the distribution of the values A and B after the selection with $A \leq B$ is non-uniform. For example,

$$p(x_i \leq B) = \frac{x_i - l_B}{u_B - l_B}$$

for $l_B \leq x_i \leq u_B$. Table 24.3 summarizes our findings about profile propagation for selections.

predicate	f'	d'	comment
$A = c$	$f'_A = f_A/d_A$	$d'_A = 1$	
$c_1 \leq A \leq c_2$	$f'_A = \frac{c_2 - c_1}{u_A - l_A} * f_A$ $f'_A = d'_A * (f_A/d_A)$	$d'_A = \frac{c_2 - c_1}{u_A - l_A} * d_A$ $d'_A = \frac{(c_2 - c_1)}{\Delta_A}$	
$A = B$	$f'_A = \frac{f}{\max(d_A, d_B)}$ $f'_A = f'_B = \frac{f_A}{n}$	$d'_A = d_A * \mathcal{Y}_{f_A/d_A}^{f_A}(f'_A)$ $d'_A = d_A * \mathcal{Y}_{f_A/d_A}^{f_A}(f'_A)$	$\Pi_A(R) \stackrel{\subseteq}{=} \Pi_B(R)$ else
$A \leq B$	$f'_A = f'_B = \frac{f_A}{2}$	$d'_A = d_A * \mathcal{Y}_{f_A/d_A}^{f_A}(f'_A)$	
$p(\mathcal{A})$	$f'_C = f_A$ $C \notin \mathcal{A} = \mathcal{F}(p)$	$d'_C = d_C * \mathcal{Y}_{d_C/f_C}^{f_A}(f'_A)$	

Table 24.3: Profile propagation for selection

Open ranges and functions There are plenty of other cases for selection predicates, which we have not discussed. Let us briefly mention a few of them. Clearly, we have:

$$\begin{aligned}
|\sigma_{A \neq c}(R)| &= |R| - |\sigma_{A=c}(R)| \\
|\sigma_{c_1 < A \leq c_2}(R)| &= |\sigma_{c_1 \leq A \leq c_2}(R)| - |\sigma_{A=c_1}(R)| \\
|\sigma_{c_1 \leq A < c_2}(R)| &= |\sigma_{c_1 \leq A \leq c_2}(R)| - |\sigma_{A=c_2}(R)| \\
|\sigma_{c_1 < A < c_2}(R)| &= |\sigma_{c_1 \leq A \leq c_2}(R)| - |\sigma_{A=c_1}(R)| - |\sigma_{A=c_2}(R)| \\
|\sigma_{c_1 \leq A}(R)| &= |\sigma_{c_1 \leq A \leq u_A}(R)| \\
|\sigma_{c_1 < A}(R)| &= |\sigma_{c_1 \leq A < u_A}(R)| \\
|\sigma_{A \leq c_2}(R)| &= |\sigma_{l_A \leq A \leq c_2}(R)| \\
|\sigma_{A < c_2}(R)| &= |\sigma_{l_A \leq A < c_2}(R)|
\end{aligned}$$

This helps to estimate the f'_A . The d'_A are left to the reader. EXC

Estimating selectivities for (user defined) functions and expressions can be done by using *computed attributes*. For example, cardinalities for selections with predicates like $g(A) = c$ for a function g can be treated by introducing an additional attribute g_A for which a profile can be established.

24.3.4 Profile Propagation for Joins

Semijoin Let us now turn to the join operator and its variants. We start with the left-semijoin and consider expressions of the type $R \triangleright \ltimes_{A=B} S$. If $\Pi_A(R) \subseteq$

$\Pi_B(S)$, then $R \bowtie_{A=B} S = R$, and no profiles change. If $\Pi_A(R) \supseteq \Pi_B(S)$, then $f'_A = f_A d_B / d_A$ and $d'_A = d_B$. If A and B are independent, we calculate

$$f'_A = \sum_{i=1}^n \frac{f_A}{d_A} p(x_i = A) p(x_i \in B) = \sum_{i=1}^n \frac{f_A}{d_A} \frac{d_A}{n} \frac{d_B}{n} = \frac{f_A d_B}{n}.$$

and

$$d'_A = \frac{d_A d_B}{n}.$$

For an attribute $C \in \mathcal{A}(R) \setminus \{A, B\}$, we have $f'_C = f'_A$ and

$$d'_C = d_C * \mathcal{Y}_{f'_C/d_C}^{f'_A}(f'_A).$$

Regular Join For the regular join $R \bowtie_{A=B} S$. Let us start with an attribute $C \in \mathcal{A}(R) \setminus \{A, B\}$. We can apply the formulas for the semijoin because $\Pi_C^D(R \bowtie_{A=B} S) = \Pi_C^D(R \bowtie_{>} S)$. For attributes $C \in \mathcal{A}(S) \setminus \{A, B\}$ remember that the join commutes.

We turn to the case where $\Pi_A(R) \subseteq \Pi_B(S)$. In this case, it is easy to see that $f'_A = f'_B = f_A f_B / d_B$ and $d'_A = d'_B = d_A$. If $\Pi_A(R) \supseteq \Pi_B(S)$, we have just to exchange the roles of R and S .

If A and B are independent, we proceed as follows. Denote again by n the domain size of A and B and the values of $D_A = \Pi_A(R)$ by $\{x_1, \dots, x_{d_A}\}$. Then, we can argue that

$$f'_A = f'_B = \sum_{i=1}^n \frac{f_A}{d_A} \frac{f_B}{d_B} p(x_i = A) p(x_i \in B) = \sum_{i=1}^n \frac{f_A}{d_A} \frac{f_B}{d_B} \frac{d_A}{n} \frac{d_B}{n} = \frac{f_A f_B}{n}.$$

Rosenthal showed that this result also holds if the condition of *fairness* holds for at least one relation [624]. A relation is called *fair* with respect to an attribute A if for the expected value $E(|\sigma_{A=x}(R)|) = |R|/n_A$ holds. In this case, the expected value for the result of the join is $(|R||S|)/n$. Note that $\Pi_A^D(R \bowtie_{A=B} S) = \Pi^D(R \bowtie_{>} S)$. Thus, we can estimate the number of distinct values as

$$d'_A = d'_B = \frac{d_A d_B}{n}.$$

Selfjoin The above formulas only apply if we are not dealing with a selfjoin. Of course, $R \bowtie_{A=B} R$ does not pose any problems. However, $R \bowtie_{A=A} R$ does, because all tuples find a join partner. The estimates are easy to derive:

$$\begin{aligned} f'_A &= \frac{f_A f_A}{d_A} \\ d'_A &= d_A \end{aligned}$$

For all attributes C other than A , $f'_C = f'_A$ and $d'_C = d_C$.

As pointed out by [24], selfjoin sizes can be used to derive an upper bound for general joins:

$$|R \bowtie_{A=B} S| \leq \frac{|R \bowtie_{A=A} R| + |S \bowtie_{B=B} S|}{2}.$$

join	f'	d'	comment
$R \bowtie_{A=B} S$	$f'_A = f_A$	$d'_A = d_A$	$\Pi_A(R) \subseteq \Pi_B(S)$
	$f'_A = \frac{f_A d_B}{d_A}$	$d'_A = d_B$	$\Pi_A(R) \supseteq \Pi_B(S)$
	$f'_A = \frac{f_A d_B}{n}$	$d'_A = \frac{d_A d_B}{n}$	else
	$f'_C = f'_A$	$d'_C = d_C * \mathcal{Y}_{d_C/f_C}^{f'_A}(f'_A)$	for $C \in \mathcal{A}(R) \setminus \{A, B\}$
$R \bowtie_{A=B} S$	$f'_A = \frac{f_A f_B}{d_B}$	$d'_A = d_A$	$\Pi_A(R) \subseteq \Pi_B(S)$
	$f'_A = \frac{f_A f_B}{n}$	$d'_A = \frac{d_A d_B}{n}$	else
$R \bowtie_{A=A} R$	$f'_A = \frac{f_A f_A}{d_A}$	$d'_A = d_A$	

Table 24.4: Profile propagation for joins

This bound, which is an immediate consequence of the Cauchy Schwarz inequality, can be used as a sanity check. Table 24.4 summarizes our findings for joins.

24.3.5 Profile Propagation for Projections

The next operator we deal with is the duplicate eliminating projection. Given our work from Chapter 4, it is not surprising, that we simply have to apply the \mathcal{D} function. For single attributes A , we have

$$f'_A = d'_A = d_A$$

For a set of attributes $\mathcal{A} = \{A_1, \dots, A_n\}$, the result cardinality of $\Pi_{\mathcal{A}}^D(R)$ is

$$\mathcal{D}\left(\prod_{i=1}^n n_{A_i}, |R|\right),$$

if n_{A_i} , the size of the domain of A_i , is known. Otherwise, we can use the estimate

$$\mathcal{D}\left(\prod_{i=1}^n d_{A_i}, |R|\right)$$

The number of distinct values in any attribute does not change, i.e. $d'_{A_i} = d_{A_i}$.

If we have a functional dependencies and $\kappa \rightarrow \mathcal{A}$ for a set of attributes \mathcal{A} and $\kappa \subset \mathcal{A}$, then

$$\Pi_{\mathcal{A}}^D(R) = \Pi_{\kappa}^D(R).$$

Further, if $|\Pi_{\mathcal{A}}^D(R)| = |R|$, we have $|\Pi_{\mathcal{A}'}^D(R)| = |R|$ for all \mathcal{A}' with $\mathcal{A}' \supseteq \mathcal{A}$.

The above estimates for the result size of a duplicate eliminating projection assumes that the attribute values are uniformly distributed, i.e., every distinct value occurs with the same probability. As we will not deal with projections any

more in this part of the book, let us complete the subject by giving an approach where each attribute value can have its own probability of occurrence. This is not unlikely, and for attributes with few possible values the following approach proposed by Yu, Zuzarte, and Sevcik is quite reasonable [?]. The assumptions are that the attributes are independent and the values of each of them are drawn by independent Bernoulli trials. Under these assumptions, they derive the following three results: a lower bound, an upper bound, and an estimate for the expected number of distinct values in the projection. In order to state these results, we need some additional notation. Let R be a relation and define $N = |R|$. Further let $G = \{A_1, \dots, A_n\}$ be a subset of the attributes of R . Define $d_i = |\Pi_{A_i}^D(R)|$ to be the number of distinct values occurring in attribute A_i . We denote values of A_i by $a_{i,1}, \dots, a_{i,d_i}$.

We wish to derive an estimate for $D_G = |\Pi_G^D(R)|$. Therefore, we model each attribute A_i by a frequency vector $f_i = (f_{i,1}, \dots, f_{i,d_i})$ where $f_{i,j}$ is the number of occurrences of the j -th distinct value $a_{i,j}$ of A_i divided by N . If, for example, A_1 has three distinct values which occur 90, 9, and 1 times in a relation with $N = 100$ elements, then f_1 becomes $(0.9, 0.09, 0.01)$.

Let us first look at bounds for D_G . Trivially, D_G is bounded from above by

$$D_G \leq \min\{N, \prod_{i=1}^n d_i\}$$

and from below by

$$D_G \geq \max_{i=1}^n d_i.$$

These bounds are very rough. This motivated Yu et al. to derive better ones.

Before we proceed, let us consider another example. Assume we have three attributes A_1, A_2 , and A_3 all with frequency vectors $f_i = (0.9, 0.09, 0.01)$ for a relation of size $N = 100$. Since we assume attribute independence, the probability of $(a_{1,3}, a_{2,3}, a_{3,3})$ is $0.01 * 0.01 * 0.01$. Thus, its occurrence in a relation of size 100 is highly unlikely. Hence, we expect D_G to be less than $27 = 3*3*3$. In general, we observe that the probability of occurrence of a tuple $(a_{1,j_1}, \dots, a_{n,j_n})$ is the product of the relative frequencies $f_{1,j_1} * \dots * f_{n,j_n}$. From this, the basic idea of the approach of Yu et al. becomes clear: we have to systematically consider all the different possibilities to multiply relative frequencies. This is nicely captured by the Kronecker product (tensor product).

Before we proceed, let us state the upper and lower bounds in case of two attributes by giving two theorems developed by Yu et al. [?].

Theorem 24.3.1 (lower bound) *For a set of attributes $\{A_1, A_2\}$ of a relation R and its frequency vectors, we define $l_{i,j}$ for $i = 1, 2$ and $1 \leq j \leq d_i$ as the minimum number of different values that have to be combined with $f_{i,j}$ given the marginals, i.e.*

$$l_{i,j} = \min\{|F| \mid F \subseteq \{1, \dots, d_{i'}\}, \forall q \notin F \ S(F) \leq f_{i,j} < S(F) + f_{i',q}\}$$

where $i' = 3 - i$ and $S(F) = \sum_{p \in F} f_{i',p}$. Further define

$$D_G^{\frac{1}{2}} = \max_{i=1,2} \sum_{j=1}^{d_i} l_{i,j}.$$

```

CalculateLowerBoundForNumberOfDistinctValues( $f_1, f_2$ )
/* frequency vectors  $f_1$  and  $f_2$  */
sort  $f_i$  ( $i = 1, 2$ ) in descending order;
for  $i = 1, 2$  {
     $i' = 3 - i$ ;
    for  $j = 1, \dots, d_i$  {
         $k = 1$ ;
        while ( $f_{i,j} > \sum_{l=1}^k f_{i',l}$ )
            ++ $k$ ;
         $l_{i,j} = k$ ;
    }
     $lb_i = \sum_{j=1}^{d_i} l_{i,j}$ ;
}
 $D_G^\perp = \max_{i=1,2} lb_i$ ;
return  $D_G^\perp$ 

```

Figure 24.2: Calculating the lower bound D_G^\perp

Then D_G^\perp is a lower bound for D_G and $D_G^\perp \geq \max_{d_1, d_2}$.

Theorem 24.3.2 (upper bound) For a set of attributes $\{A_1, A_2\}$ of a relation R and its frequency vectors, we define $u_{i,j}$ for $i = 1, 2$ and $1 \leq j \leq d_i$ as the maximum number of different values that can be combined with $f_{i,j}$ given the marginals, i.e.

$$u_{i,j} = \min\{d_{i'}, f_{i,j}N\}$$

where $i' = 3 - i$. Further define

$$D_G^\top = \min_{i=1,2} \sum_{j=1}^{d_i} u_{i,j}.$$

Then D_G^\top is a lower bound for D_G and $D_G^\top \leq \min\{N, \prod_{i=1}^n d_i\}$.

The algorithm in Fig. 24.2 calculates the lower bound D_G^\perp . Calculating the upper bound D_G^\top is much easier. For each $f_{i,j}$, we compute $u_{i,j}$ by simply comparing $f_{i,j}N$ and $d_{i'}$. Adding up the $u_{i,j}$ for each attribute and taking the lesser of the two sums gives the desired result.

Let us start by repeating the definition of the Kronecker product of two matrices $A = (a_{i,j})$ and $B = (b_{i,j})$ of dimension $n \times m$ and $n' \times m'$. The result $A \otimes B$ is a matrix of dimension $nn' \times mm'$. The general definition is

$$A \otimes B = \begin{pmatrix} a_{1,1}B & a_{1,2}B & \dots & a_{1,m}B \\ a_{2,1}B & a_{2,2}B & \dots & a_{2,m}B \\ \dots & \dots & \dots & \dots \\ a_{n,1}B & a_{n,2}B & \dots & a_{n,m}B \end{pmatrix}.$$

The estimate can not be calculated easily. First, we calculate the Kronecker product $f_G = f_1 \otimes \dots \otimes f_n$ of all frequency vectors. Note that to every value

combination $v \in \Pi_{A_1}^D(R) \times \dots \times \Pi_{A_n}^D(R)$ there corresponds exactly one component in f_G , which contains its probability of occurrence. With this observation, it is easy to derive the following theorem, in which we denote by $f_{G,i}$ the i -th component of f_G and by M its length, i.e. $M = \prod_{i=1}^n d_i$. Further remember that $N = |R|$.

Theorem 24.3.3 (estimate) *Let the following assumptions hold:*

1. *The data distributions of individual attributes in G are independent.*
2. *For the value combinations v_i , its occurrence is the result of an independent Bernoulli trial, with the success (occurrence) probability $f_{G,i}$.*
3. *The occurrences of individual possible value combinations are independent.*

Then, the expected number of distinct values D_G is

$$E[D_G] = M - \sum_{i=1}^M (1 - f_{G,i})^N.$$

The algorithm for computing the estimate is given in Fig. 24.3. In the first, most expensive phase, it constructs the Kronecker product. Then, the simple calculations according to the theorem follow. A more efficient implementation would calculate the Kronecker product only implicitly. Further, the frequency vectors may not be completely known but only a part of it via some histogram. As was also shown by Yu et al., end-biased histograms (coming soon) are optimal under the following error metrics. Let $\hat{D}_{G,\text{hist}}$ be the estimate derived for a histogram. The error function they consider is

$$\mathcal{E}_{\text{abs}} = |\hat{D}_G - \hat{D}_{G,\text{hist}}|.$$

24.3.6 Profile Propagation for Division

As a starting point, we use an observation made by Merrett and Otoo [?]. Assume we are given two sets X and Y , which are both subsets of a finite domain D with $|D| = n$ elements. Then $|X| < |Y|$ implies that $X \not\subseteq Y$. Otherwise, we can calculate the probability of $X \supseteq Y$ as

$$p(X \supseteq Y) = \binom{|X|}{|Y|} / \binom{n}{|Y|}$$

Now let R and S be two relations with $\mathcal{A}(R) = \{A, B\}$ and $\mathcal{A}(S) = \{B\}$. A value $a \in \Pi_A^D(R)$ is contained in the result of $R \div B$ if and only if $\Pi_B(\sigma_{A=a}(R)) \supseteq S$. Hence, for any such a , $\bar{f}_A = f_A/d_A$ and n_B equal to the size of the common domain of $R.B$ and $S.B$, we can calculate the survival probability as

$$\binom{\bar{f}_A}{|S|} / \binom{n_B}{|S|}$$

```

EstimateNumberOfDistinctValues( $f_1, \dots, f_n$ )
  /* frequency vectors  $f_i$  */
  /* step 1: calculate  $f_G = f_1 \otimes \dots \otimes f_n$  */
   $f_G = f_1$ ;
  for ( $i = 2; i \leq n; ++i$ ) {
     $f^{\text{old}} = f_G$ ;
     $f_G = \epsilon$ ; // empty vector
    for ( $j = 1; j \leq |f^{\text{old}}|; ++j$ ) {
      for ( $k = 1; k \leq d_i; ++k$ ) {
         $f_G = \text{push\_back}(f_G, f_j^{\text{old}} \times f_{i,j})$ ; // append a value to a vector
      }
    }
  }
  /* step 2: compute the expected number of distinct value combinations */
   $S = 0$ ;
  for ( $j = 1, j \leq M; ++j$ ) { //  $M = \text{length}(f_G)$ 
     $S += (1 - f_j)^N$ ;
  }
   $\hat{D}_G = M - S$ ;
  return  $\hat{D}_G$ ;

```

Figure 24.3: Calculating the estimate for D_G

provided that $\overline{f_A} \geq |S|$ and R is a set. Denote by f'_A and d'_A the cumulated frequency and the number of distinct values for attribute A in the result of $R \div S$. Then we have the estimate

$$f'_A = d'_A = d_A * \frac{\binom{\overline{f_A}}{|S|}}{\binom{n_B}{|S|}}$$

in case R is a set.

If R is a bag, we must be prepared to see duplicates in $\sigma_{A=a}(R)$. In this case we can adjust the above formula to

$$f'_A = d'_A = d_A * \frac{\binom{x_A}{|S|}}{\binom{n}{|S|}}$$

where $x_A = \mathcal{D}(x_A, n_A)$, and n_A is the size of the domain of $R.A$.

If there is some variance among the number of distinct values associated with the $a \in \Pi_A^D(R)$, the estimate will be rough. To cure this, we need better information. Define for each $a \in \Pi_A^D(R)$ the number h_a to be the number of distinct b values occurring for it, i.e. $h_a = |\Pi_B^D(\sigma_{A=a}(R))|$. Then we could estimate f'_A and d'_A as follows:

$$f'_A = d'_A = \sum_{a \in \Pi_A^D(R)} \frac{\binom{h_a}{|S|}}{\binom{n}{|S|}}$$

Name	Definition	Error minimized
median(\tilde{x})	$\begin{cases} x_{(n+1)/2} & n \text{ odd} \\ (x_{n/2} + x_{n/2+1})/2 & n \text{ even} \end{cases}$	$\mathcal{E}_1 = \sum_{i=1}^n x_i - \hat{x} $
mean(\bar{x})	$1/n \sum_{i=1}^n x_i$	$\mathcal{E}_2 = \sqrt{\sum_{i=1}^n (x_i - \hat{x})^2}$
middle	$(\max(x) + \min(x))/2$	$\mathcal{E}_\infty = \max_{i=1}^n x_i - \hat{x} $
q-value	$\sqrt{\max(X) \min(X)}$	$\mathcal{E}_q = \max_{i=1}^n \max\{x_i/\hat{x}, \hat{x}/x_i\}$

Table 24.5: Approximations of a set of numbers by a single number

Keeping h_a for every possible a may not be practical. However, if the number of distinct values in $H = \{h_a | a \in \Pi_A^D(R)\}$ is small, we can keep the number of distinct a values for each possible h_a . Assume $H = \{h_1, \dots, h_k\}$ and define

$$g_i = |\{a \in \Pi_A^D(R) | h_a = h_i\}|,$$

then we have the estimate

$$f'_A = d'_A = \sum_{i=0, h_i \geq |S|}^k g_i \binom{h_i}{|S|} / \binom{n}{|S|}.$$

24.3.7 Remarks

NULL Values Our profile is not really complete for attributes which can have NULL values. To deal with these, we need to extend our profiles by the frequency d_A^\perp with which NULL occurs in an attribute A of some relation. It is straightforward to extend the above profile to deal with this additional count.

Uniformity is not sufficient As we have seen, even if all attributes are uniformly distributed, which is rarely the case in practice, the result of algebraic operators may no longer be uniformly distributed. As a consequence, we need to be concerned with the approximation of the true distribution of values.

Sets of Attributes Note that nothing prevents us to use the formulas developed above for selections and joins if A and B are attribute sets instead of single attributes. We just have to know or calculate d_A for sets of attributes \mathcal{A} .

24.4 Approximation of a Set of Values

24.4.1 Approximations and Error Metrics

Assume we have a set of values $x = \{x_1, \dots, x_n\}$. The task we want to tackle is to approximate this set of values by a single value. The left two columns

of Table 24.5 show the names and definitions of some possible approximations. Whereas mean and median are well known, the other two may be not. The middle is defined as the value exactly between the minimum and maximum of X . Hence, the distance from the middle to either extreme is the same. The q-value needs some further restriction: the values in X must be larger than zero. For our purposes, this restriction is not bad since execution costs are typically larger than zero and frequencies are mostly larger than zero if they are not exactly zero. The latter case needs some special attention if we use something like the q-value, which we could also term *geometric* or *multiplicative middle*.

Let us take a look at a simple example. Assume $X = \{1, 2, 9\}$. Then we can easily calculate the approximations summarized in the following table:

median	mean	middle	q-value
2	4	5	3

Which of these approximations is the best one? The answer depends on the error function we wish to minimize. Therefore, the rightmost column of Table 24.5 shows some error functions, which are minimized by the approximation defined in the same line. The variable \hat{x} denotes the estimate whose error is to be calculated. For \mathcal{E}_2 there exist plenty of equivalent formulations, where we think of two error measures as being equivalent, if and only if they result in the same minimum. Some important alternatives are $1/n \sum (x_i - \hat{x})^2$, $1/(n-1) \sum (x_i - \hat{x})^2$ (empirical variance), and simply $\sum (x_i - \hat{x})^2$.

A nice property half of the approximations give us are error bounds. These are \mathcal{E}_∞ and \mathcal{E}_q . Define the *spread* s of x as $\max(x) - \min(x)$. Then, given the middle m of x , we have for every $x_i \in x$ that

$$m - s/2 \leq x_i \leq m + s/2.$$

Thus, we have a symmetric, additive error bound for all elements in x . Define the *geometric spread* as $s = \sqrt{\max(x)/\min(x)}$. Then we have a symmetric, multiplicative error bound for all elements x_i in x given by

$$(1/s)q \leq x_i \leq sq$$

if q is the geometric middle. The following table shows the possible errors for all approximations of our example set $X = \{1, 2, 9\}$:

	median	mean	middle	geo. mean
	2	4	5	3
\mathcal{E}_1	8	10	11	9
\mathcal{E}_2	7.1	6.2	6.4	6.4
\mathcal{E}_∞	7	5	4	6
\mathcal{E}_q	4.5	4	5	3

Which of these error metrics and, hence, which approximation is the best? Obviously, this depends on the application. In the query compiler context, \mathcal{E}_1 plays no role that we are aware of. \mathcal{E}_2 plays a predominant role as it is used to approximate values in a given histogram bucket. This has not come by sharp reasoning about the best possibility but merely by the existence of a huge body of literature in this area. Currently, the other two error metrics, \mathcal{E}_∞ and \mathcal{E}_q , play minor roles. But this will change.

24.4.2 Example Applications

Let us discuss some example applications relevant to building a query compiler. Assume we have to come up with the execution time (CPU usage) for some function. This could be a simple arithmetic function built into our system, a hash function executed for a hash-join, the CPU time used to dereference a TID if the according page is in memory, the CPU time needed to search a B-tree page residing in the buffer, or the CPU time needed to load a page from secondary storage into the buffer. Careful as we are, we measure the function's execution time several times. Almost certainly, the numbers will not be same for every execution, except maybe for the simplest functions. To come up within a single number, we need to approximate the set of numbers derived from our measurements. If the function is going to be executed many times within a query execution plan (in which it occurs), we need to cost the average case and the mean is the approximation of choice. We will see other applications when we discuss histograms.

24.5 Approximation with Linear Models

24.5.1 Linear Models

In this section, we want to approximate a given set of points (x_i, y_i) ($1 \leq i \leq m$) by a linear combination \hat{f} of given functions Φ_j , $1 \leq j \leq n$. The general assumption is that $m > n$. We define the estimation function \hat{f} as

$$\hat{f}(x) := \sum_{j=1}^n c_j \Phi_j(x)$$

for coefficients $c_j \in \mathbb{R}$. The estimates \hat{y}_i for y_i are then derived from \hat{f} by

$$\hat{y}_i := \hat{f}(x_i) = \sum_{j=1}^n c_j \Phi_j(x_i).$$

Note that the functions Φ_j are not necessarily linear functions. For example, we could use polynomials $\Phi_j(x) = x^{j-1}$. Further, there is no need for x to be a single number. It could as well be a vector \vec{x} .

It is convenient to state our approximation problem in terms of vectors and matrices. Let (x_i, y_i) be the points we want to approximate and Φ_j , $1 \leq j \leq n$ be some functions. We define the *design matrix* $A \in \mathbb{R}^{m \times n}$, $A = (a_{i,j})$ by

$$a_{i,j} = \Phi_j(x_i)$$

or, equivalently, as an explicit matrix

$$A = \begin{pmatrix} \Phi_1(x_1) & \Phi_2(x_1) & \Phi_3(x_1) & \dots & \Phi_n(x_1) \\ \Phi_1(x_2) & \Phi_2(x_2) & \Phi_3(x_2) & \dots & \Phi_n(x_2) \\ & & \dots & & \\ \Phi_1(x_m) & \Phi_2(x_m) & \Phi_3(x_m) & \dots & \Phi_n(x_m) \end{pmatrix} \quad (24.5)$$

Assume we wish to approximate the points by a polynomial of degree $n - 1$. Then, $\Phi_i(x) = x_{i-1}$ and the design matrix becomes

$$A = \begin{pmatrix} 1 & (x_1)^1 & (x_1)^2 & \dots & (x_1)^{n-1} \\ 1 & (x_2)^1 & (x_2)^2 & \dots & (x_2)^{n-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & (x_m)^1 & (x_m)^2 & \dots & (x_m)^{n-1} \end{pmatrix}.$$

In the simplest case, where we want to use a linear function $\hat{f}(x) = c_1 + c_2x$ to approximate the points, the design matrix becomes

$$A = \begin{pmatrix} 1 & x_1 \\ 1 & x_2 \\ \vdots & \vdots \\ 1 & x_m \end{pmatrix}.$$

As an example consider the three points

$$(1, 20), (2, 10), (3, 60).$$

The design matrix becomes

$$A = \begin{pmatrix} 1 & 1 \\ 1 & 2 \\ 1 & 3 \end{pmatrix} \quad (24.6)$$

For every column vector $\vec{c} = (c_1, c_2)^T$

$$A\vec{c}$$

gives the result of \hat{f} for all points. Clearly, \vec{c} should be determined such that the deviation of $A\vec{c}$ from $\vec{y} = (y_1, \dots, y_m)^T$ becomes minimal.

The deviation could be zero, that is $A\vec{c} = \vec{y}$. However, remember our assumption that $m > n$. This means that we have more equations than variables. Thus, we have an overdetermined system of equations and it is quite unlikely that a solution to this system of equations exists. This motivates our goal to find an approximation as good as possible. Next, we formalize this goal.

Often used measures for deviations or distances of two vectors are based on norms.

Definition 24.5.1 (norm) *Let S be a linear space. Then a function $\|x\| : S \rightarrow \mathbb{R}$ is called a norm if and only if it has the following three properties:*

1. $\|x\| > 0$ unless $x = 0$
2. $\|\lambda x\| = |\lambda| \|x\|$
3. $\|x + y\| \leq \|x\| + \|y\|$

Various norms, called p norms can be found in the literature. Let $x \in \mathbb{R}^n$ and $p \geq 1$ where $p = \infty$ is possible. Then

$$\|x\|_p = \left(\sum_{i=1}^n |x_i|^p \right)^{\frac{1}{p}}.$$

The most important norms are the l_1 , l_2 , and l_∞ norms:

$$\begin{aligned} \|x\|_1 &= |x_1| + \dots + |x_n| \\ \|x\|_2 &= \sqrt{(x_1)^2 + \dots + (x_n)^2} \\ \|x\|_\infty &= \max_{i=1}^n |x_i| \end{aligned}$$

Using these norms, we can define distance functions d_1 , d_2 , and d_∞ . For two vectors x and y in \mathbb{R}^n , we define

$$\begin{aligned} d_1(x, y) &= \|x - y\|_1 \\ d_2(x, y) &= \|x - y\|_2 \\ d_\infty(x, y) &= \|x - y\|_\infty \end{aligned}$$

It should be clear, that these define the error measures \mathcal{E}_1 , \mathcal{E}_2 , and \mathcal{E}_∞ , which we used in Sec. 24.4. The only missing error function is \mathcal{E}_q . We immediately fill this gap, and start with the one dimensional case.

Definition 24.5.2 (Q-paranorm in R) Define for $x \in R$,

$$\|x\|_Q = \begin{cases} \infty & \text{if } x \leq 0 \\ 1/x & \text{if } 0 < x \leq 1 \\ x & \text{if } 1 \leq x \end{cases}$$

$\|\cdot\|_Q$ is called Q-paranorm.

Note that for $x > 0$, $\|x\|_Q = \max(x, 1/x)$. The multivariate case is a straightforward extension using the maximum over all components:

Definition 24.5.3 (Q-paranorm in R^n) For $x \in R^n$, define

$$\|x\|_Q = \max_{i=1}^n \|x_i\|_Q.$$

We denote this paranorm by l_q .

Definition 24.5.4 (paranorm) Let S be a linear space. Then a function $\|x\| : S \rightarrow R$ is called a paranorm if and only if the following two properties hold:

1. $\|x\| \geq 0$
2. $\|x + y\| \leq \|x\| + \|y\|$

The Q-paranorm is a norm, hence the name. The only missing part is the distance function stated next. Let x and y be two vectors in \mathbb{R}^n , where $y = (y_1, \dots, y_n)^T$ with $y_i > 0$. Then we define

$$d_q(x, y) = \|x/y\|_Q$$

where we define x/y for two column vectors $x, y \in \mathbb{R}^n$ as follows:

$$x/y = (x_1/y_1, \dots, x_n/y_n)^T.$$

Between norms there exist some inequalities. For all vectors $x \in \mathbb{R}^n$, we have

$$\begin{aligned} \|x\|_2 &\leq \|x\|_1 \leq \sqrt{n}\|x\|_2 \\ \|x\|_\infty &\leq \|x\|_2 \leq \sqrt{n}\|x\|_\infty \\ \|x\|_\infty &\leq \|x\|_1 \leq n\|x\|_\infty \end{aligned}$$

For l_q , no such inequality exists as $\|x\|_q$ approaches infinity as x approaches zero.

We can now formally state the approximation problem. Let $A \in \mathbb{R}^{m \times n}$ be the design matrix and (x_i, y_i) , $1 \leq i \leq m$ be a set of points, and $\vec{y} = (y_1, \dots, y_m)$. The goal is to find a vector $\vec{a}^* \in \mathbb{R}^n$ minimizing $d(A\vec{a}, \vec{y})$. That is, we look for $\vec{a}^* \in \mathbb{R}^n$ such that

$$d(A\vec{a}^*, \vec{y}) = \min_{\vec{a} \in \mathbb{R}^n} d(A\vec{a}, \vec{y}) \tag{24.7}$$

\vec{a}^* is then called *solution* of the approximation problem or *best approximation*.

For different l (d), we get different problems. For l_1 the problem is called *quantile regression*. We will not deal with it here, since we do not know of any application of it in the database context. The solutions for the problems for l_2 , l_∞ , and l_q are discussed in subsequent sections, after we have given some example applications of what needs to be approximated in a DBMS. Before we proceed, let us give the solutions for approximating the points $(1, 20)$, $(2, 10)$, $(3, 60)$ with a linear function $\alpha + \beta x$. The following table shows the values of x , y and estimates for y produced by the best approximations \hat{f}_{l_2} , \hat{f}_{l_∞} , \hat{f}_{l_q} , which minimize l_2 , l_∞ , and l_q , resp. Additionally, we give the α and β of the best approximations as well as their quality measured by l_1 , l_2 and l_q .

x	y	\hat{f}_{l_2}	\hat{f}_{l_∞}	\hat{f}_{l_q}
1	20	10	5	10
2	10	30	25	20
3	60	50	45	30
α		20	20	10
β		-10	-15	0
l_2		14.1421	15	19.1485
l_∞		20	15	30
l_q		3	4	2

Let us repeat some general insights into approximation problems as defined above. Thereby, we follow the exposition of Watson [?]. We start with stating theorems on the existence of a solution. The following two theorems only apply to norms. That is, they do not apply to l_q . However, as we will see later, solutions under l_q exist.

Theorem 24.5.5 (Existence 1) *Let M denote a compact set in a normed linear space. Then to each point g of the space there exists a point of M closest to g .*

Compactness is a sufficient but not necessary condition.

Theorem 24.5.6 (Existence 2) *Let M be a finite dimensional subspace of a normed linear space S . Then there exists a best approximation in M to any point of S .*

The next point to consider is the uniqueness of a solution. Proving the uniqueness of a solution is easy, if the norm is *strictly convex*.

Definition 24.5.7 ((strictly) convex) *Let $f(x)$ be a function on the elements x of a linear space S . Then $f(x)$ is convex if*

$$f(\lambda x_1 + (1 - \lambda)x_2) \leq \lambda f(x_1) + (1 - \lambda)f(x_2)$$

for all $x_1, x_2 \in S$ and $0 \leq \lambda \leq 1$.

If $0 < \lambda < 1$ implies strict inequality in the above inequality, $f(x)$ is called strictly convex.

It is easy to show that all l_p norms for $p \neq \infty$ are strictly convex and that l_∞ and l_q are convex, but not strictly convex. For strictly convex norms, it is easy to show that a solution is unique.

Theorem 24.5.8 *In a strictly convex normed linear space S , a finite dimensional subspace M contains a unique best approximation to any point of S .*

Although l_∞ and l_q are not strictly convex, under certain circumstances a unique best approximation exists for them. This is discussed in subsequent sections.

Considering the above, one might conjecture that l_2 approximation is much simpler than l_∞ or l_q approximation. This is indeed the case. We will not repeat all the findings from approximation theory and the algorithms developed. There are plenty of excellent textbooks on this matter. We highly recommend the excellent book of Golub and van Loan [?], which discusses l_2 approximation and several algorithms to solve them (e.g. QR factorization and SVD). Other good books to refresh one's knowledge on matrix algebra are [?, ?, ?, ?]. Überhuber wrote another good book discussing l_2 approximation, QR factorization and SVD [?]. In the context of statistics, many different regression models exist to approximate a given set of data. An excellent overview is provided by [?]. Another good reading, not only in this context, is the book by Press, Teukolsky, Vetterling, and Flannery [597]. Before reading these books, it might be helpful to repeat some linear algebra and some basics of matrices. An excellent book for doing so was written by Schmidt and Trenkler [?]. The only book we know of that discusses approximation under l_∞ , is the one by Watson, already cited above [?]. Approximation under l_q is not discussed in any textbook. Hence, we must refer to the original article [?, ?]. In any case, since mathematics is quite involved at times, we give explicit algorithms only for the approximation by a linear function. For all other cases, we refer to the literature.

ToDo

24.5.2 Example Applications

In this section, we give some examples of approximation problems occurring in the database context. As we will see, different problems demand different norms. Additionally, we sketch how to use approximations. The details are left to the reader as an exercise.

Disk seek times

There exist small benchmarks, which measure the disk seek time for travelling n cylinders (see Sec. 4.1). To cover for random errors, many measurements are taken. The task is to find the parameters d and c_0 , $1 \leq i \leq 4$, for the disk seek time formula from Sec. 4.1:

$$seektime(d) = \begin{cases} c_1 + c_2\sqrt{d} & d \leq c_0 \\ c_3 + c_4d & d > c_0 \end{cases}$$

Since many seeks occur during the processing of a single query, l_2 is the appropriate norm. On the surface, we seem to have a problem using a simple linear model. However, we can approximate the parts $c_1 + c_2\sqrt{d}$ and $c_3 + c_4d$ for several distinct c_0 either by trying a full range of values for c_0 or by a binary search. The solution for c_0 we then favor is the one in which the maximum of the errors on both parts becomes minimal. A second problem is the occurrence of \sqrt{d} , since this does not look linear. However, choosing $\Phi_1 = 1$ and $\Phi_2(x) = \sqrt{x}$ will work fine.

Another method is to transform a set of points (x_i, y_i) with two (injective) transformation functions t_x and t_y into the set of points $(t_x(x_i), t_y(y_i))$. Then this set is approximated and the result is transformed back. While using this approach, special attention has to be paid to the norm, as it can change due to the transformation. We see examples of this later on in Sec. 24.5.6.

Functions sensitive to parameter size

Another example is to approximate the execution time of a hash function on string values. As its calculation depends on the length of the input string, measurements can be taken for various lengths. Using l_2 as a norm is perfect, because the hash function is typically executed many times during a hash join or hash group [?].

Approximation of frequency densities and distributions

We start by demonstrating the usage of approximating functions for cardinality estimation. Then, we look at the choice of error metrics for estimating selectivity results and the influence of cardinality estimation errors on joins.

Let R be a relation and A one of its attributes. Let (x_i, f_i) denote the frequency f_i with which the value x_i occurs in attribute A . Typically, only those values x_i are written down and approximated for which $f_i \neq 0$. We further assume that the x_i are sorted, i.e. $x_i < x_{i+1}$. Using the methods to come, we can approximate this set of points by a function $\hat{f}(x)$. To calculate

the output cardinality of a selection $\sigma_{A=c}(R)$, we can simply return $\hat{f}(c)$ as an estimate. Hence it is a good choice to use l_q (see below for strong arguments).

To calculate the result cardinality for a range query of the form $\sigma_{c_1 \leq A \leq c_2}(R)$, we distinguish several cases. First, assume that the domain of A is discrete and the number of values between c_1 and c_2 is small. Then, we can calculate the result quite efficiently by

$$\sum_{c_1 \leq x \leq c_2} \hat{f}(x).$$

If the number of values between c_1 and c_2 is too large for an explicit summation, we can apply speed-up techniques if the function \hat{f} has a simple form. For example, if \hat{f} is a linear function $\hat{f}(x) = \alpha + \beta x$, the above sum can be calculated very efficiently. If the number of values between c_1 and c_2 is very large and no efficient form for the above sum can be found, or if we do not have a discrete domain, we can use the integral to approximate the sum. Thus, we use the right-hand side of

$$\sum_{c_1 \leq x \leq c_2} \hat{f}(x) \approx \int_{c_1}^{c_2} \hat{f}(x) dx$$

to approximate the sum by evaluating an expression which is hopefully less expensive to evaluate.

Yet another solution is the following. Instead of approximating (x_i, f_i) directly, we approximate (x_i, g_i) where $g_i = \sum_{j \leq i} (f_j)$. Let us denote the approximation of this *cumulated frequency distribution* by \hat{g} . Then the result cardinality of a range query of the form $\sigma_{c_1 \leq A \leq c_2}(R)$ can be simply calculated by

$$\hat{g}(c_2) - \hat{g}(c_1).$$

Minimizing error propagation. Let us assume that the purpose of our approximation is to estimate the output cardinalities of selections on relations R_i , i.e. $\sigma_{p_i}(R_i)$ for $i = 1, \dots, n$. The results of these cardinality estimations are then used to find the optimal order of subsequent joins. More specifically, assume we have to find the optimal query execution plan for the following expression:

$$\sigma_{p_1}(R_1) \bowtie \dots \bowtie \sigma_{p_n}(R_n), \quad (24.8)$$

where we intentionally left out all the join predicates. Ioanidis and Christodoulakis pointed out that errors propagate exponentially through joins [384]. Denote by s_i the cardinality of $\sigma_{p_i}(R_i)$ and by \hat{s}_i its estimate. Further assume that independence holds. This means, that s_i can be written as $f_i |R_i|$, where f_i is the selectivity of p_i . Denote by $f_{i,j}$ the selectivity of the join predicate between R_i and R_j , if it exists. Otherwise we define $f_{i,j} = 1$. The result of joining a subset $x \subseteq \{R_1, \dots, R_n\}$ has cardinality

$$s_x = \left(\prod_{R_i \in x} f_i \right) \left(\prod_{R_i, R_j \in x} f_{i,j} \right) \left(\prod_{R_i \in x} |R_i| \right)$$

Denote by \hat{f}_i the estimate for the selectivities of the p_i and assume that the join selectivities have been estimated correctly (which, of course, is difficult in

EXC

practice). Then the estimated cardinality of the result of joining the relations in x is

$$\begin{aligned}
\hat{s}_x &= \left(\prod_{R_i \in x} \hat{f}_i \right) \left(\prod_{R_i, R_j \in x} f_{i,j} \right) \left(\prod_{R_i \in x} |R_i| \right) \\
&= \left(\prod_{R_i \in x} f_i / f_i \right) \left(\prod_{R_i \in x} \hat{f}_i \right) \left(\prod_{R_i, R_j \in x} f_{i,j} \right) \left(\prod_{R_i \in x} |R_i| \right) \\
&= \left(\prod_{R_i \in x} \hat{f}_i / f_i \right) \left(\prod_{R_i \in x} f_i \right) \left(\prod_{R_i, R_j \in x} f_{i,j} \right) \left(\prod_{R_i \in x} |R_i| \right) \\
&= \left(\prod_{R_i \in x} \hat{f}_i / f_i \right) s_x
\end{aligned}$$

where some i belong to the category with $\hat{f}_i / f_i < 1$ and others to the one with $\hat{f}_i / f_i > 1$. Remember that during dynamic programming, all subsets of relations are considered. Especially those subsets occur in which all relations belong to one category only. Hence, building on the cancellation of errors by mixing them from different categories is not a true option. Instead, we should minimize

$$\prod_{R_i \in x} \max\{f_i / \hat{f}_i, \hat{f}_i / f_i\}$$

in order to minimize errors and error propagation. This product can be minimized by minimizing each of its factors. This means that if we want to minimize error propagation, we have to minimize the multiplicative error \mathcal{E}_q for estimating the cardinalities of selections based on equality. This finding can obviously be generalized to any kind of selections. Thus, for cardinality estimations for selections (and joins, or cardinality estimation in general) the q-error is the error and the geometric middle the approximation of choice.

Error bounds guaranteeing plan optimality. Let us give another strong argument for minimizing the multiplicative error \mathcal{E}_q . Let us consider again the join expression given in 24.8. Further, denote by f_i the correct selectivity of $\sigma_{A_i=c_i}$ and by \hat{f}_i some estimate. If the plan generator uses the correct cardinalities, it produces the optimal plan. Given the estimates, it might produce another plan. The question is, how far can the cardinality estimates deviate from the true cardinalities and the optimal plan remains the same. More formally, denote by P the optimal plan under the correct cardinalities f and by \hat{P} the optimal plan under the estimates \hat{f} . Then, we can restate the above question to whether there exists a condition on \hat{f} such that if this condition holds then $\hat{P} = P$. The nice truth is that such conditions exist and they involve the Q paranorm.

In the simplest case, let us assume that the expression given in 24.8 is used to evaluate a star query under an ASI cost function without considering cross products. From Sec. 3.2.2 we can conclude that the optimal join order for star queries starts with the center relation and orders the satellite relations according to their rank. The rank of a relation R is defined as $\text{rank}(R_i) = (T(R_i) - 1) / C(R_i)$, where $C(S)$ are some fixed per tuple costs and $T(R_i) =$

$f_{0,i}f_i|R_i|$, if $f_{0,i}$ is the join selectivity of the join of R_i with the center relation R_0 . Thus, $\hat{P} = P$ if f and \hat{f} result in the same ordering of the relations. Since $f(x) = (x - 1)/c$ is monotonically increasing for constants c , we can conclude that the ordering is indeed the same as long as for all $i \neq j$ we have

$$f_{0,i}f_i|R_i| < f_{0,j}f_j|R_j| \iff f_{0,i}\hat{f}_i|R_i| < f_{0,j}\hat{f}_j|R_j|$$

which is equivalent to

$$\frac{f_i r_i}{f_j r_j} < 1 \iff \frac{\hat{f}_i r_i}{\hat{f}_j r_j} < 1$$

for $r_i = f_{0,i}|R_i|$. We now show that if

$$\left\| \frac{f_i}{\hat{f}_i} \right\|_Q < \min_{i \neq j} \sqrt{\left\| \frac{f_i r_i}{f_j r_j} \right\|_Q}$$

for all i , then $P = \hat{P}$. This condition implies the much weaker condition that for all $i \neq j$

$$\left\| \frac{\hat{f}_i}{f_i} \right\|_Q \left\| \frac{\hat{f}_j}{f_j} \right\|_Q < \left\| \frac{f_i r_i}{f_j r_j} \right\|_Q \quad (24.9)$$

To show the claim, it suffices to show that $(f_i r_i)/(f_j r_j) < 1$ implies $(\hat{f}_i r_i)/(\hat{f}_j r_j) < 1$. This follows from

$$\begin{aligned} \frac{\hat{f}_i r_i}{\hat{f}_j r_j} &= \frac{\hat{f}_i f_j f_i r_i}{f_i \hat{f}_j f_j r_j} \\ &= \left(\frac{\hat{f}_i f_j}{f_i \hat{f}_j} \right) / \left(\left\| \frac{f_i r_i}{f_j r_j} \right\|_Q \right) \quad (*) \\ &\leq \left(\left\| \frac{\hat{f}_i}{f_i} \right\|_Q \left\| \frac{f_j}{\hat{f}_j} \right\|_Q \right) / \left(\left\| \frac{f_i r_i}{f_j r_j} \right\|_Q \right) \\ &< 1 \end{aligned}$$

where (*) follows from $(f_i r_i)/(f_j r_j) < 1$.

For chain queries,

$$\left\| \frac{\hat{f}_i}{f_i} \right\|_Q < \min_{i \neq j-1} \sqrt{\left\| \frac{f_i f_{i,i+1} |R_i|}{f_j f_{j,j-1} |R_j|} \right\|_Q}$$

implies $\hat{P} = P$. For tree queries, things are a little more complex. In the following, we assume that $R_{i'}$ is a relation connected to R_i , which we denote by $R_{i'} - R_i$. Then

$$\left\| \frac{\hat{f}_i}{f_i} \right\|_Q < \min_{i \neq j, R_{i'} - R_i, R_{j'} - R_j} \sqrt{\left\| \frac{f_i f_{i,i'} |R_i|}{f_j f_{j,j'} |R_j|} \right\|_Q}$$

implies $\hat{P} = P$.




Figure 24.4: Q-error and plan optimality

From cardinality estimation error bounds to cost error bounds.

Positions of interesting values

Consider an integer attribute A of a relation R with values in $[l, u]$. Imagine a situation where almost all values in $[l, u]$ occur with some frequency in A . Only few holes exist. Denote by W the set of values in $[l, u]$ that do not occur in A . Again, the selection is of the form $\sigma_{c_1 \leq A \leq c_2}(R)$. In the case of holes, the above summation for range queries leads to a wrong result. It would be better to calculate the result cardinality as in

$$\sum_{c_1 \leq x \leq c_2} \hat{f}(x) - \sum_{c_1 \leq x \leq c_2, x \in W} \hat{f}(x)$$

For this to work, we have to know where the holes in $[l, u]$ are. If there are only a few of them, we can memorize them. If they are too many to be stored, we can approximate them as follows. Let $W = \{w_1, \dots, w_m\}$. Then, we can use approximation techniques to approximate the set of points (i, w_i) , $1 \leq i \leq m$. Depending on the size of the interval $[l, u]$ either l_∞ or l_q is the appropriate norm. Similarly, the *peaks*, i.e. the distinct values occurring in the attribute A of R , can be approximated if there are only a few of them in the domain of A .

Number of distinct values

24.5.3 Linear Models Under l_2

Now that we know that the solution to our problem exists and is unique, we continue by characterizing it. Let S be a linear space (say \mathbb{R}^2) and $s \in S$ some point. Further denote by G some linear subspace of S (say a straight line). For any $g \in G$, we can define the residual vector $g - f$. Using residuals, we can characterize the unique solution quite easily and intuitively. Exactly the vector $g^* \in G$ is closest to f whose residual $f - g$ is orthogonal to G . Remember that two vectors are orthogonal if and only if their scalar product is zero. Now we can characterize the solution to our approximation problem under l_2 .

Theorem 24.5.9 (Characterization) *Let S be a linear space and G a subspace. An element g^* is the best approximation of a point $s \in S$ if and only if*

$$\langle g^* - f, g \rangle = 0$$

holds. That is, if the error is orthogonal to all elements in G .

Since we are used to solve equations for x , we rewrite our problem to $A\vec{x} = b$. That is, the vector \vec{x} replaces the coefficient vector c . Using Theorem 24.5.9, we must have that $A\vec{x}^* - b$ is orthogonal to the range of A . The range of a matrix $A \in \mathbb{R}^{m \times n}$ is defined as $\mathcal{R}(A) = \{Ax | x \in \mathbb{R}^n\}$. Let a_i be i -th column vector of A and $\vec{x} = (x_1, \dots, x_n)^T$. Then, the best approximation can be found by solving the following system of linear equations, which is called (*Gauß*) *normal equations*:

$$\begin{array}{cccccc} \langle a_1, a_1 \rangle x_1 & + & \langle a_2, a_1 \rangle x_2 & + & \dots & + & \langle a_n, a_1 \rangle x_n & = & \langle b, a_1 \rangle \\ \langle a_1, a_2 \rangle x_1 & + & \langle a_2, a_2 \rangle x_2 & + & \dots & + & \langle a_n, a_2 \rangle x_n & = & \langle b, a_2 \rangle \\ \vdots & & \vdots & & \vdots & & \vdots & & \vdots \\ \langle a_1, a_n \rangle x_1 & + & \langle a_2, a_n \rangle x_2 & + & \dots & + & \langle a_n, a_n \rangle x_n & = & \langle b, a_n \rangle \end{array}$$

or, using matrix notation, we get

$$A^T A \vec{x} = A^T \vec{b} \quad (24.10)$$

This system of linear equations can be solved by many different approaches. Some fast and numerically stable approaches are QR decomposition and singular value decomposition (SVD). Both leave the conditioning of the problem unchanged. QR decomposition can only be applied if the the matrix has full rank (see below). Otherwise, one has to keep up with variants of QR decomposition or SVD. Hence, we will briefly discuss SVD. We will not give any algorithms. The interested reader is referred to [?].

Before we proceed with SVD, let us repeat some basics on matrices. A special matrix is the identity matrix $I \in \mathbb{R}^{n \times n}$ with $I = (\delta_{i,j})_{i,j}$, $1 \leq i \leq n$. Matrices can have plenty of properties. Here are some of them.

Definition 24.5.10 (rank) *The rank of a matrix A , denoted by $\text{rank}(A)$, is the rank of the subspace $\mathcal{R}(A)$.*

Definition 24.5.11 (full rank) *A matrix $A \in \mathbb{R}^{m \times n}$, $m > n$ has full rank if its rank is n .*

Definition 24.5.12 (symmetric) *A matrix $A \in \mathbb{R}^n$ is symmetric if and only if $A^T = A$.*

Note that for all matrices $A \in \mathbb{R}^{m \times n}$, we always have that AA^T and $A^T A$ are symmetric.

Definition 24.5.13 (idempotent) *A matrix $A \in \mathbb{R}^{n \times n}$ is idempotent if and only if $AA = A$.*

Definition 24.5.14 (inverse) A Matrix $A^{-1} \in \mathbb{R}^{n \times n}$ is the inverse of a matrix $A \in \mathbb{R}^{n \times n}$ if and only if $A^{-1}A = AA^{-1} = I$.

A matrix for which the uniquely determined inverse exists is called *regular*.

Definition 24.5.15 (orthogonal) A matrix $A \in \mathbb{R}^{n \times n}$ is orthogonal if and only if $AA^T = A^T A = I$.

Let us use a simple, operational, recursive definition of the determinant.

Definition 24.5.16 (determinant) Let $A \in \mathbb{R}^{n \times n}$ be a matrix. We define the determinant of A as $\det(A) = a_{1,1}$ if $n = 1$. Otherwise, we define

$$\det(A) = \sum_{j=1}^n (-1)^{i+j} a_{i,j} \det(A_{i,j})$$

where $A_{i,j} \in \mathbb{R}^{(n-1) \times (n-1)}$ results from A by eliminating the i -th row and the j -th column.

Definition 24.5.17 (characteristic polynomial) Let $A \in \mathbb{R}^{n \times n}$ be a matrix. The characteristic polynomial is defined as

$$P_n(z; A) := \det(A - zI) = \begin{vmatrix} (a_{1,1} - z) & a_{1,2} & \cdots & a_{1,n} \\ a_{2,1} & (a_{2,2} - z) & \cdots & a_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n,1} & a_{n,2} & \cdots & (a_{n,n} - z) \end{vmatrix}$$

Definition 24.5.18 (Eigenvalue) Let $A \in \mathbb{R}^{n \times n}$ be a matrix and $P_n(z; A)$ its characteristic polynomial. Any root λ_i of $P_n(z; A)$, i.e. $P(z; A)(\lambda_i) = 0$ is called Eigenvalue. The set of Eigenvalues is denoted by

$$\lambda(A) := \{\lambda_1, \dots, \lambda_k\}$$

and is called spectrum of A .

Definition 24.5.19 (similar) Two matrices $A, B \in \mathbb{R}^{n \times n}$ are similar if and only if there exists a regular matrix $X \in \mathbb{R}^{n \times n}$ such that $B = X^{-1}AX$.

Two similar matrices have the same Eigenvalues, as can be seen from the following theorem.

Theorem 24.5.20 Let $A, B \in \mathbb{R}^{n \times n}$ be two similar matrices. Then they have the same characteristic polynomial.

Definition 24.5.21 (generalized inverse) A matrix $A^- \in \mathbb{R}^{n \times m}$ is the generalized inverse, or g-inverse, of a matrix $A \in \mathbb{R}^{m \times n}$, if $AA^-A = A$ holds.

Every matrix and, hence, every vector has a g-inverse. For regular matrices, the g-inverse and the inverse coincide. In general, the g-inverse is not uniquely determined. Adding some additional properties makes it unique.

Definition 24.5.22 (Moore-Penrose inverse) A matrix $A^+ \in \mathbb{R}^{n \times m}$ is the Moore-Penrose inverse of a matrix $A \in \mathbb{R}^{m \times n}$ if the following conditions hold:

1. $AA^+A = A$
2. $A^+AA^+ = A^+$
3. $(A^+A)^T = A^+A$
4. $(AA^+)^T = AA^+$

For every matrix and, hence, every vector there exists a uniquely determined Moore-Penrose inverse. In case A is regular, $A^+ = A^{-1}$ holds. If A is symmetric, then $A^+A = AA^+$. If A is symmetric and idempotent, then $A^+ = A$. Further, all of A^+A , AA^+ , $I - A^+A$, and $I - AA^+$ are idempotent. Here are some equalities holding for the Moore-Penrose inverse:

$$(A^+)^+ = A \quad (24.11)$$

$$(A^T)^+ = (A^+)^T \quad (24.12)$$

$$(A^T A)^+ = A^+ (A^T)^+ \quad (24.13)$$

$$(AA^T)^+ = (A^T)^+ A^+ \quad (24.14)$$

$$A^T AA^+ = A^T \quad (24.15)$$

$$A^+ AA^T = A^+ \quad (24.16)$$

The following theorem states the existence of a decomposition of any matrix into regular/orthogonal submatrices.

Theorem 24.5.23 (singular value decomposition) Let $A \in \mathbb{R}^{m \times n}$ be a matrix. Then there exist an orthogonal matrix $U \in \mathbb{R}^{m \times m}$ and an orthogonal matrix $V \in \mathbb{R}^{n \times n}$ such that

$$U^T AV = S$$

such that $S \in \mathbb{R}^{m \times n}$ is of the form

$$S = \text{diag}(s_1, \dots, s_k)$$

with $k = \min(m, n)$ and, further

$$s_1 \geq s_2 \geq \dots \geq s_r > s_{r+1} = \dots = s_k = 0$$

holds where $r = \text{rank}(A)$.

For a proof and algorithms to calculate the SVD of an arbitrary matrix see the book by Golub and Loan [?]. Another proof can be found in the book by Harville [?]. The diagonal elements s_i of S , which is orthogonal equivalent to A , are called *singular values*. From

$$S^T S = (U^T AV)^T (U^T AV) = V^T A^T U U^T AV = V^{-1} A^T AV$$

it follows that $S^T S$ and $A^T A$ are similar. Since $S^T S = \text{diag}(s_1^2, \dots, s_r^2, 0, \dots, 0)$ and similar matrices have the same spectrum, it follows that

$$s_i = \sqrt{\lambda_i}$$

for $\lambda_i \in \lambda(A^T A)$, $1 \leq i \leq n$.

Define $S^{-1} = \text{diag}(1/s_1, \dots, 1/s_r, 0, \dots, 0)$ and $A^+ = VS^{-1}U^T$. From

$$\begin{aligned} AA^+A &= (USV^T)(VS^{-1}U^T)(USV^T) \\ &= USS^{-1}SV^T \\ &= USV^T \\ &= A \end{aligned}$$

and

$$\begin{aligned} A^+AA^+ &= (VS^{-1}U^T)(USV^T)(VS^{-1}U^T) \\ &= VS^{-1}SS^{-1}U^T \\ &= VS^{-1}U^T \\ &= A^+ \end{aligned}$$

we see that $A^+ = VS^{-1}U^T$ is a g-inverse of A . The reader is advised to check the remaining conditions of the Moore-Penrose inverse.

Remember that we have to solve $A^T A \vec{x} = A^T \vec{b}$ for \vec{x} in order to find the best approximation for our set of data points. Set $\vec{x} = A^+ \vec{b}$. Then

$$\begin{aligned} A^T A \vec{x} &= A^T A A^+ \vec{b} \\ &= A^T \vec{b} \end{aligned}$$

where we used Eqn. 24.15. Hence, the Moore-Penrose inverse² solves our problem. Moreover, the solution can be obtained easily from the singular value decomposition.

Approximation by a linear function

Assume we are given m points (x_i, y_i) , $1 \leq i \leq m$ and wish to approximate them by a linear function $f(x) = \alpha + \beta x$. The design matrix, b and \vec{x} then are

$$A = \begin{pmatrix} 1 & x_1 \\ 1 & x_2 \\ \vdots & \vdots \\ 1 & x_m \end{pmatrix} \quad x = \begin{pmatrix} \alpha \\ \beta \end{pmatrix}, \quad b = \begin{pmatrix} y_1 \\ \vdots \\ y_m \end{pmatrix}$$

The resulting system of normal equations

$$\begin{pmatrix} m & \sum_{i=1}^m x_i \\ \sum_{i=1}^m x_i & \sum_{i=1}^m (x_i)^2 \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \begin{pmatrix} \sum_{i=1}^m y_i \\ \sum_{i=1}^m x_i y_i \end{pmatrix}$$

²The Greville algorithm to calculate the Moore-Penrose inverse directly is described in [?].

has the solution

$$\alpha = \frac{\sum_{i=1}^m (x_i)^2 \sum_{i=1}^m y_i - \sum_{i=1}^m x_i y_i \sum_{i=1}^m x_i}{m \sum_{i=1}^m (x_i)^2 - (\sum_{i=1}^m x_i)^2}, \beta = \frac{m \sum_{i=1}^m x_i y_i - \sum_{i=1}^m x_i \sum_{i=1}^m y_i}{m \sum_{i=1}^m (x_i)^2 - (\sum_{i=1}^m x_i)^2}$$

Note that this is a very nice formula as new points arrive or are deleted, only the sums have to be updated and the quotients to be calculated. There is no need to look at the other points again.

24.5.4 Linear Models Under l_∞

Let $A \in \mathbb{R}^{m \times n}$ be a matrix, where $m > n$, and $b \in \mathbb{R}^m$ a vector. The problem we solve in this section is to

$$\text{find } \vec{a} \in \mathbb{R}^n \text{ to minimize } \|r(\vec{a})\|_\infty \quad (24.17)$$

where

$$r(\vec{a}) = \vec{b} - A\vec{a} \quad (24.18)$$

The components of the vector $r(\vec{a})$ are denoted by $r_i(\vec{a})$.

As pointed out earlier, l_∞ is a convex norm. Hence, a solution exists. Since l_∞ is not strictly convex, the uniqueness of the solution is not guaranteed. To solve problem 24.17 by following the approach proposed by Watson [?]. We start by characterizing the solution, continue with the conditions under which uniqueness holds, make some more observations, and finally derive an algorithm for the case $n = 2$, i.e. we find a best approximation by a linear function. Although only few applications in databases exist for l_∞ , it is very useful to find a best approximation under l_q if we want to approximate by a function $e^{\beta + \alpha x}$ (see Sec. 24.5.6).

Assume we have a best solution \vec{a} . Then, for some indices i , $r_i(\vec{a})$ attains the maximum, i.e. $r_i(\vec{a}) = \|r(\vec{a})\|_\infty$. Otherwise, a better solution would exist. We denote the set of indices where the maximum is attained by $\bar{I}(\vec{a})$. We further denote by $\theta_i(\vec{a})$ the sign of $r_i(\vec{a})$. Thus $r_i(\vec{a}) = \theta_i(\vec{a}) \|r(\vec{a})\|_\infty$ for all $i \in \bar{I}$. The following theorem gives a characterization of the solution.

Theorem 24.5.24 *A vector $\vec{a} \in \mathbb{R}^n$ solves problem 24.17 if and only if there exists a subset I of \bar{I} with $|I| \leq n + 1$ and a vector $\vec{\lambda} \in \mathbb{R}^m$ such that*

1. $\lambda_i = 0$ for all $i \notin I$,
2. $\lambda_i \theta_i \geq 0$ for all $i \in I$, and
3. $A^T \vec{\lambda} = \vec{0}$.

The set I in the theorem is called an *extremal subset* of a solution \vec{a} .

There are two important corollaries to this theorem.

Corollary 24.5.25 *Let \vec{a} solve problem 24.17. Then \vec{a} solves an l_∞ approximation problem in \mathbb{R}^{n+1} obtained by restricting the components of $r(\vec{a})$ to some particular $n + 1$ components. If A has rank t , then the components of $r(\vec{a})$ may be restricted to a particular $t + 1$ components.*

Corollary 24.5.26 *Let \vec{a} solve problem 24.17 and let I be chosen according to Theorem 24.5.24 such that $\lambda_i \neq 0$ for all $i \in I$. Further let \vec{d} be another solution to 24.17. Then*

$$r_i(\vec{d}) = r_i(\vec{a}).$$

Hence, not surprisingly, any two solutions have the same residuals for components where the maximum is attained. The theorem and its first corollary state that we need at most $t+1$ solutions for a matrix A of rank t . The next theorem shows that at least $t+1$ indices exist where the maximum is attained.

Theorem 24.5.27 *If A has rank t , a solution to problem 24.17 exists for which $|\bar{I}| \geq t+1$.*

Thus, any submatrix of A consisting of a subset of the rows of A , which correspond to the indices contained in $\bar{I}(\vec{a})$, must have rank t for some solution \vec{a} to problem 24.17.

The above theorems and corollaries indicate that the clue to uniqueness is the rank of subsets of rows of A . The following definition captures this intuition.

Definition 24.5.28 (Haar condition) *A matrix $R \in Rmn$, where $m \geq n$ satisfies the Haar condition if and only if every submatrix consisting of n rows of A is nonsingular.*

Finally, we can derive uniqueness for those A which satisfy the Haar condition:

Theorem 24.5.29 *If A satisfies the Haar condition, the solution to problem 24.17 is unique.*

Obviously, we need to know, whether the Haar condition holds for a matrix A . Remember that we want to approximate a set of points by a linear combination of functions Φ_j , $1 \leq j \leq n$. From the points (x_i, y_i) , $1 \leq i \leq m$, and the Φ_j , the design matrix A is derived as shown in Equation 24.5. If the Φ_j form a Chebyshev set, the design matrix will fulfill the Haar condition.

Definition 24.5.30 (Chebyshev set) *Let X be a closed interval of R . A set of continuous function $\Phi_1(x), \dots, \Phi_n(x)$, $\Phi_i : X \rightarrow R$, is called a Chebyshev set, if every non-trivial linear combination of these functions has at most $n-1$ zeros in X .*

Assume the x_i are ordered, that is $x_i < x_{i+1}$ for $1 \leq i < m$. Further, it is well-known that the set of polynomials $\Phi_j = x^{j-1}$, $1 \leq j \leq n$, forms a Chebyshev set on any interval X . From now on, we assume that our x_i are ordered, that is $x_1 < \dots < x_m$. Further, we define $X = [x_1, x_m]$. We also assume that the matrix A of Problem 24.17 is defined as given in Equation 24.5, where the Φ_j are continuous functions from X to \mathbb{R} .

We still need some more knowledge in order to build an algorithm. The next definition will help to derive a solution for subsets I of $\{1, \dots, m\}$ with $|I| = n+1$.

Definition 24.5.31 (alternating set) Let \vec{a} be a vector in \mathbb{R}^n . We say that $r(\vec{a})$ alternates s times, if there exists points $x_{i_1}, \dots, x_{i_s} \in \{x_1, \dots, x_m\}$ such that

$$r_{i_k}(\vec{a}) = -r_{i_{k+1}}(\vec{a})$$

for $1 \leq k < s$. The set $\{x_{i_1}, \dots, x_{i_s}\}$ is called an alternating set for \vec{a} .

Theorem 24.5.32 Let (x_i, y_i) , $1 \leq i \leq m$, be an ordered set of points with $x_i \leq x_{i+1}$ for $1 \leq i < m$. Define $X = [x_1, x_m]$. Further let Φ_j , $1 \leq j \leq n$ be a Chebyshev set on X . Define $A = (a_{i,j})$, where $1 \leq i \leq m$, $1 \leq j \leq n$, and $a_{i,j} = \Phi_j(x_i)$. Then, a vector $\vec{a} \in \mathbb{R}^n$ solves Problem 24.17 if and only if there exists an alternating set with $n + 1$ points for a .

Consider again the example where we want to approximate the three points $(1, 20)$, $(2, 10)$, and $(3, 60)$ by a linear function. We saw that the solution to our problem is $\hat{f}_\infty(x) = -15 + 20x$. The following table gives the points, the value of \hat{f}_∞ , the residuals, including their signs.

x	y	\hat{f}_∞	r_i
1	20	5	+15
2	10	25	-15
3	60	45	+15

As Theorem 24.5.32 predicts, the signs of the residual alternate.

The proof of Lemma 24.5.32 uses the following lemma (see [?]).

Lemma 24.5.33 Let (x_i, y_i) , $1 \leq i \leq m$, be an ordered set of points with $x_i \leq x_{i+1}$ for $1 \leq i < m$. Define $X = [x_1, x_m]$. Further let Φ_j , $1 \leq j \leq n$ be a Chebyshev set on X . Define the $n \times n$ determinant

$$\Delta_i = \Delta(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_{n+1})$$

where

$$\Delta(x_1, \dots, x_n) = \det \begin{vmatrix} \Phi_1(x_1) & \dots & \Phi_n(x_1) \\ \vdots & \vdots & \vdots \\ \Phi_1(x_n) & \dots & \Phi_n(x_n) \end{vmatrix} \quad (24.19)$$

Then

$$\text{sign}(\Delta_i) = \text{sign}(\Delta_{i+1}), \quad \forall 1 \leq i \leq n.$$

Let us take a closer look at Theorem 24.5.32 in the special case where $m = 3$, i.e. we have exactly three points (x_{i_1}, y_{i_1}) , (x_{i_2}, y_{i_2}) , and (x_{i_3}, y_{i_3}) . We find the best linear approximation $\hat{f}(x) = \alpha + \beta x$ under l_∞ by solving the following equations:

$$\begin{aligned} y_{i_1} - (\alpha + \beta x_{i_1}) &= -1 * \lambda \\ y_{i_2} - (\alpha + \beta x_{i_2}) &= +1 * \lambda \\ y_{i_3} - (\alpha + \beta x_{i_3}) &= -1 * \lambda \end{aligned}$$

where λ represents the value of $\|r(\vec{a})\|_\infty$ for the solution \vec{a} to be found. Solving these equations results in

$$\begin{aligned}\lambda &= \frac{y_{i_2} - y_{i_1}}{2} - \frac{(y_{i_3} - y_{i_1})(x_{i_2} - x_{i_1})}{2(x_{i_3} - x_{i_1})} \\ \beta &= \frac{y_{i_2} - y_{i_1}}{x_{i_2} - y_{i_1}} - \frac{2\lambda}{x_{i_2} - x_{i_1}} \\ \alpha &= y_{i_1} + \lambda - x_{i_1}\beta\end{aligned}$$

The algorithm starts with three arbitrary points with indices i_1, i_2 , and i_3 with $x_{i_1} < x_{i_2} < x_{i_3}$. Next, it derives α, β , and λ using the solutions to our equations 24.20-24.20. Then, the algorithm tries find new indices j_1, j_2, j_3 by exchanging one of the i_j with some k such that λ will be increased. Obviously, we use a k that maximizes the deviation from the best approximation \hat{f} of i_1, i_2, i_3 , i.e.

$$\|y_k - \hat{f}(x_k)\|_\infty = \max_{i=1,\dots,m} \|y_i - f(x_i)\|_\infty.$$

Depending on the position of x_k in the sequence i_1, i_2, i_3 and the sign of the residual we determine the i_j to be exchanged with k .

- $x_k < x_{i_1}$
if $(\text{sign}(y_k - \hat{f}_k) == \text{sign}(y_{i_1} - \hat{f}_{i_1}))$
then $j_1 = k, j_2 = i_2, j_3 = i_3$
else $j_1 = k, j_2 = i_1, j_3 = i_2$
- $x_{i_1} < x_k < x_{i_2}$
if $(\text{sign}(y_k - \hat{f}_k) == \text{sign}(y_{i_1} - \hat{f}_{i_1}))$
then $j_1 = k, j_2 = i_2, j_3 = i_3$
else $j_1 = i_1, j_2 = k, j_3 = i_2$
- $x_{i_2} < x_k < x_{i_3}$
if $(\text{sign}(y_k - \hat{f}_k) == \text{sign}(y_{i_2} - \hat{f}_{i_2}))$
then $j_1 = i_1, j_2 = k, j_3 = i_2$
else $j_1 = i_1, j_2 = i_2, j_3 = k$
- $x_k > x_{i_3}$
if $(\text{sign}(y_k - \hat{f}_k) == \text{sign}(y_{i_3} - \hat{f}_{i_3}))$
then $j_1 = i_1, j_2 = i_2, j_3 = k$
else $j_1 = i_2, j_2 = i_3, j_3 = k$

The above rules are called *exchange rules*. In general, they state that if k falls between two indices, the one with the same sign as r_k is replaced by k . If k is smaller than the smallest index (larger than the largest index), we consider two cases. If the smaller (largest) index has the same sign of its residue as k , we exchange it with k ; otherwise we exchange it with the largest (smallest) index. Stated this way, we can use the exchange rules for cases where $n > 2$.

Algorithm 24.5 summarizes the above considerations.

In case $n > 2$, the above algorithm remains applicable. We just have to use the general exchange rule and provide a routine solving the following system of equations for x_i and λ :

BestLinearApproximationUnderChebyshevNorm

1. Choose arbitrary i_1, i_2, i_3 with $x_{i_1} < x_{i_2} < x_{i_3}$.
(e.g. equi-distant i_j .)
2. Calculate the solution for the system of equations.
This gives us an approximation function $\hat{f}(x) = \alpha + \beta x$ and λ .
3. Find an x_k for which the deviation of \hat{f} from the given data is maximized.
Call this maximal deviation λ_{\max} .
4. If $\lambda_{\max} - \lambda > \epsilon$ for some small ϵ
then apply the exchange rule using x_k and go to step 2.
(The ϵ is mainly needed for rounding problems with floating point numbers.)
5. Return α, β, λ .

Figure 24.5: Algorithm for best linear approximation under l_∞

$$\begin{aligned}
 a_{1,1}x_1 + a_{1,2}x_2 + a_{1,n}x_n &= -\lambda \\
 a_{2,1}x_1 + a_{2,2}x_2 + a_{2,n}x_n &= +\lambda \\
 &\dots \quad \dots \\
 a_{n+1,1}x_1 + a_{n+1,2}x_2 + a_{n+1,n}x_n &= (-1)^{n+1}\lambda
 \end{aligned}$$

24.5.5 Linear Models Under l_q

Let (x_i, y_i) for $1 \leq i \leq m$ be a set of points with $y_i > 0$, which we again want to approximate by a linear combination of a given set of functions Φ_j , $1 \leq j \leq n$. This time, we measure the deviation by applying l_q . That is, we want to find a coefficients a_j such that the function

$$\hat{f}(x) = \sum_{j=1}^n a_j \Phi_j(x)$$

minimizes

$$\max_{i=1, \dots, m} \max \left\{ \frac{y_i}{\hat{f}(x_i)}, \frac{\hat{f}(x_i)}{y_i} \right\}.$$

Let \vec{a} and \vec{b} be two vectors in \mathbb{R}^n with $b_i > 0$. Then, we define $\vec{a}/\vec{b} = \frac{\vec{a}}{\vec{b}} = (a_1/b_1, \dots, a_n/b_n)^t T$.

Let $A \in \mathbb{R}^{m \times n}$ be a matrix, where $m > n$ and $\vec{b} = (b_1, \dots, b_m)^t T$ be a vector in \mathbb{R}^m with $b_i > 0$. Then we can state the problem as

$$\text{find } \vec{a} \in \mathbb{R}^n \text{ that minimizes } \|A\vec{a}/\vec{b}\|_Q \quad (24.20)$$

under the constraint that $\alpha_i^t T > 0$, $1 \leq i \leq m$, for all row vectors α_i of A .

Alternatively, we can modify A by “dividing” it by \vec{b} . We need some notations to do so. Let $\vec{b} = (b_1, \dots, b_m)^t T$ be a vector in \mathbb{R}^m . Define $\text{diag}(\vec{b})$ to be the $m \times m$ diagonal matrix which contains the b_i in its diagonal and is zero outside the diagonal. For vectors \vec{b} with $b_i > 0$, we can define $\vec{b}^{-1} = (1/b_1, \dots, 1/b_m)^t T$.

Using these notations, we can define

$$A' = \text{diag}(\vec{b}^{-1})A$$

In the special case of univariate polynomial approximation with $\hat{f}(x) = a_1 + a_2x + \dots + a_nx^{n-1}$ the matrix A' has the form

$$A' = \begin{pmatrix} 1/y_1 & x_1/y_1 & \dots & x_1^{n-1}/y_1 \\ 1/y_2 & x_2/y_2 & \dots & x_2^{n-1}/y_2 \\ \vdots & \vdots & \dots & \vdots \\ 1/y_m & x_m/y_m & \dots & x_m^{n-1}/y_m \end{pmatrix}. \tag{24.21}$$

Keeping the trick with A' in mind, it is easy to see that Problem 24.20 can be solved, if we can solve the general problem

$$\text{find } \vec{a} \in \mathbb{R}^n \text{ that minimizes } \|A\vec{a}\|_Q. \tag{24.22}$$

The following proposition ensures that a solution to this general problem exists. Further, since $\|A\vec{a}\|_Q$ is convex, the minimum is a global one.

Proposition 24.5.1 *Let $A \in \mathbb{R}^{m,n}$ such that $\mathcal{R}(A) \cap \mathbb{R}_{>0}^m \neq \emptyset$. Then $\|A \cdot \|_Q$ attains its minimum.*

Recall that l_q is subadditive and convex. Further it is lower semi-continuous (see also [?, p. 52]). However, it is not strictly convex. Hence, as with l_∞ , we expect uniqueness to hold only under certain conditions.

We need some more notation. Let $A \in \mathbb{R}^{m,n}$. We denote by $\mathcal{R}(A) = \{A\vec{a} \mid \vec{a} \in \mathbb{R}^n\}$ the *range* of A and by $\mathcal{N}(A) = \{\vec{a} \in \mathbb{R}^n \mid A\vec{a} = 0\}$ the *nullspace* of A .

Problem (24.22) can be rewritten as the following constrained minimization problem:

$$\min_{(\vec{a}, q) \in \mathbb{R}^n \times \mathbb{R}} q \quad \text{subject to} \quad \frac{1}{q} \leq A\vec{a} \leq q \quad \text{and} \quad q \geq 1. \tag{24.23}$$

The Lagrangian of (24.23) is given by

$$L(\vec{a}, q, \lambda^+, \lambda^-, \mu) := q - (\lambda^+)^T(q - A\vec{a}) - (\lambda^-)^T(A\vec{a} - \frac{1}{q}) - \mu(q - 1).$$

Assume that $\mathcal{R}(A) \cap \mathbb{R}_{>0}^m \neq \emptyset$. Then the set $\{(\vec{a}, q) : \frac{1}{q} \leq A\vec{a} \leq q \text{ and } q \geq 1\}$ is non-empty and closed and there exists (\vec{a}, q) for which we have strong inequality in all conditions. Then the following *Karush-Kuhn-Tucker* conditions are necessary and sufficient for $(\hat{\vec{a}}, \hat{q})$ to be a minimizer of (24.23), see, e.g., [?, p. 62]: there exist $\hat{\lambda}^+, \hat{\lambda}^- \in \mathbb{R}_{\geq 0}^m$ and $\hat{\mu} \geq 0$ such that

$$\nabla_{\vec{a}} L(\hat{\vec{a}}, \hat{q}, \hat{\lambda}^+, \hat{\lambda}^-, \hat{\mu}) = A^T \lambda^+ - A^T \lambda^- = 0 \tag{24.24}$$

$$\frac{\partial}{\partial q} L(\hat{\vec{a}}, \hat{q}, \hat{\lambda}^+, \hat{\lambda}^-, \hat{\mu}) = 1 - \sum_{i=1}^m \hat{\lambda}_i^+ - \frac{1}{\hat{q}^2} \sum_{i=1}^m \hat{\lambda}_i^- - \hat{\mu} = 0 \tag{24.25}$$

and for $i = 1, \dots, m$,

$$\hat{\lambda}_i^+ \left(\hat{a} - (\hat{A}\hat{a})_i \right) = 0, \quad (24.26)$$

$$\hat{\lambda}_i^- \left((\hat{A}\hat{a})_i - \frac{1}{\hat{q}} \right) = 0, \quad (24.27)$$

$$\hat{\mu}(\hat{q} - 1) = 0.$$

Assume that $1_m \notin \mathcal{R}(A)$, where 1_m is the vector with all components 1. Then $\hat{q} > 1$ and consequently $\hat{\mu} = 0$. Furthermore, it is clear that not both $\hat{\lambda}_i^+$ and $\hat{\lambda}_i^-$ can be positive because the conditions $\hat{q} = (A\hat{a})_i$ and $\frac{1}{\hat{q}} = (A\hat{a})_i$ cannot be fulfilled at the same time, since $\hat{q} > 1$.

Setting $\hat{\lambda} := \hat{\lambda}^+ - \hat{\lambda}^-$, we can summarize our findings (24.24) - (24.27) in the following theorem.

Theorem 24.5.34 *Let $A \in \mathbb{R}^{m,n}$ such that $\mathcal{R}(A) \cap \mathbb{R}_{>0}^m \neq \emptyset$ and $1_m \notin \mathcal{R}(A)$. Then (\hat{a}, \hat{q}) solves (24.23) if and only if there exists $\hat{\lambda} \in \mathbb{R}^m$ such that*

i) $A^T \hat{\lambda} = 0$.

ii) $q = q \sum_{\hat{\lambda}_i > 0} \hat{\lambda}_i + \frac{1}{q} \sum_{\hat{\lambda}_i < 0} \hat{\lambda}_i$.

iii) $\hat{\lambda}_i = 0$ if $\frac{1}{\hat{q}} < (A\hat{a})_i < q$.

iv) if $\hat{\lambda}_i > 0$ then $(A\hat{a})_i = \hat{q}$ and if $\hat{\lambda}_i < 0$ then $(A\hat{a})_i = 1/\hat{q}$.

Remark. We see that $1 < \hat{q} = (A\hat{a})_i$ implies $\text{sign} \left((A\hat{a})_i - 1 \right) = 1$ and that $1 > 1/\hat{q} = (A\hat{a})_i$ implies $\text{sign} \left((A\hat{a})_i - 1 \right) = -1$; whence $\hat{\lambda}_i \left((A\hat{a})_i - 1 \right) \geq 0$. For our approximation problem (24.20) this means that the residuum $\hat{f}(x_i) - b_i$ fulfills $\hat{\lambda}_i (\hat{f}(x_i) - b_i) \geq 0$.

Under certain conditions, problem (24.22) has a unique solution which can be simply characterized. Let us start with some straightforward considerations in this direction. If $\mathcal{N}(A) \neq \{\vec{0}\}$, then we have for any minimizer \hat{a} of $\|A \cdot\|_Q$ that $\hat{a} + \beta$, $\beta \in \mathcal{N}(A)$ is also a minimizer. In particular, we have that $\mathcal{N}(A) \neq \{\vec{0}\}$ if

- $m < n$,
- $m \geq n$ and A is not of full range, i.e., $\text{rank}(A) < n$.

In these cases, we cannot have a unique minimizer. Note further, that if $1_m \in \mathcal{R}(A)$, then the minimum of $\|A \cdot\|_Q$ is 1 and the set of minimizers is given by

$$A^+ 1_m + \mathcal{N}(A),$$

where A^+ denotes the Moore-Penrose inverse of A . Of course, this can easily be checked using the methods of Sec. 24.5.3.

In the following, we restrict our attention to the case $m > n$ and $\text{rank}(A) = n$.

The following proposition considers $(n + 1, n)$ -matrices.

Proposition 24.5.2 *Let $A \in \mathbb{R}^{n+1,n}$ such that $\mathcal{R}(A) \cap \mathbb{R}_{>0}^{n+1} \neq \emptyset$, $1_m \notin \mathcal{R}(A)$ and $\text{rank}(A) = n$. Then $\|A \cdot\|_Q$ has a unique minimizer if and only if the Lagrange multipliers $\hat{\lambda}_i$, $i = 1, \dots, n + 1$ are not zero.*

By $\text{spark}(A)$ we denote the smallest number of rows of A which are linearly dependent. In other words, any $\text{spark}(A) - 1$ rows of A are linearly independent. For the 'spark' notation we also refer to [?].

Examples. 1. We obtain for the matrix

$$A := \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 1 \end{pmatrix}, \quad \text{rg}(A) = 3, \text{spark}(A) = 3.$$

2. The matrix (m, n) -matrix A in (24.21) is the product of the diagonal matrix $\text{diag}(1/b_i)_{i=1}^m$ with positive diagonal entries and a Vandermonde matrix. Hence, it can easily be seen that $\text{spark}(A) = n + 1$. If an (m, n) -matrix A has $\text{spark}(A) = n + 1$, then A fulfills the Haar condition.

Proposition 24.5.2 can be reformulated as follows:

Corollary 24.5.35 *Let $A \in \mathbb{R}^{n+1,n}$ such that $\mathcal{R}(A) \cap \mathbb{R}_{>0}^{n+1} \neq \emptyset$ and $1_m \notin \mathcal{R}(A)$. Then $\|A \cdot\|_Q$ has a unique minimizer if and only if $\text{spark}(A) = n + 1$.*

The result can be generalized by the following theorem.

Theorem 24.5.36 *Let $A \in \mathbb{R}^{m,n}$ such that $\mathcal{R}(A) \cap \mathbb{R}_{>0}^m \neq \emptyset$. Suppose that $\text{spark}(A) = n + 1$. Then $\|A \cdot\|_Q$ has a unique minimizer which is determined by $n + 1$ rows of A , i.e., there exists an index set $J \subset \{1, \dots, m\}$ of cardinality $|J| = n + 1$ such that $\|A \cdot\|_Q$ and $\|A|_J \cdot\|_Q$ have the same minimum and the same minimizer. Here $A|_J$ denotes the restriction of A to the rows which are contained in the index set J . We call such index set J an extremal set.*

Of course the condition $\text{spark}(A) = n + 1$ is not necessary for $\|A \cdot\|_Q$ to have a unique minimizer as the following example shows.

Example. The matrices

$$A := \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ -1 & \frac{1}{2} \\ -4 & 2 \end{pmatrix}, \quad \text{and} \quad A := \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ -4 & 4 \\ -1 & 1 \end{pmatrix}$$

have both $\text{spark}(A) = 2$. By some following considerations, we obtain for both problems that the minimum of $\|A \cdot\|_Q$ is $\hat{q} = 2$. However, in the first problem the minimizer is uniquely determined by $\hat{a} = (\frac{1}{2}, 2)^T$ while the whole line $c(\frac{1}{2}, 1)^T + (1 - c)(\frac{3}{2}, 2)^T$, $c \in [0, 1]$ minimizes the functional in the second case. For $(\frac{1}{2}, 1)^T$ we have $\text{sign}(\hat{\lambda}_1, \hat{\lambda}_2, \hat{\lambda}_3, \hat{\lambda}_4) = (-1, 0, 1, -1)$ while the pattern is $(0, 1, 1, -1)$ for $(\frac{3}{2}, 2)^T$ and $(0, 0, 1, -1)$ within the line bounded by these points.

By Theorem 24.5.36, a method for finding the minimizer of $\|A \cdot\|_Q$ would be to compute the unique minimizers of the $\binom{m}{n+1}$ subproblems $\|A|_J \cdot\|_Q$ for

all index sets J of cardinality $n + 1$ and to take the largest minimum \hat{a} and the corresponding \hat{a} as minimizer of the original problem. For our line problem there exist $\binom{m}{3} = \mathcal{O}(m^3)$ of these subproblems. In the following section, we give another algorithm which is also based on Theorem 24.5.36, but ensures that the value a enlarges for each new choice of the subset J . Since there is only a finite number of such subsets we must reach a stage where no further increase is possible and J is an extremal set. In normed spaces such methods are known as *ascent methods*, see [?].

In this section, we suggest a detailed algorithm for minimizing $\|A \cdot\|_Q$, where we restrict our attention to the line problem

$$\max_{i=1,\dots,m} \max \left\{ \frac{b_i}{\beta + \alpha x_i}, \frac{\beta + \alpha x_i}{b_i} \right\}. \quad (24.28)$$

i.e., to the matrix A in (24.21) with $n = 2$.

Corollary 24.5.37 *Let (x_i, b_i) , $i = 1, 2, 3$ be given points with pairwise distinct $x_i \in \mathbb{R}$ and positive b_i , $i = 1, 2, 3$. Then the minimum \hat{q} and the minimizer $\hat{a} \in \mathbb{R}^2$ of (24.28) are given by $\hat{q} = \|\hat{q}_1\|_Q$ and*

$$\begin{pmatrix} \hat{\beta} \\ \hat{\alpha} \end{pmatrix} = \frac{1}{x_2 - x_1} \begin{pmatrix} x_2 & -x_1 \\ -1 & 1 \end{pmatrix} \begin{pmatrix} b_1 \hat{q}_1 \\ b_2 \hat{q}_2 \end{pmatrix},$$

where

$$\hat{q}_1 := \begin{cases} \sqrt{\frac{r_2}{1-r_1}} & \text{if } r_1 < 0 \text{ and } r_2 > 0, \\ \sqrt{\frac{1-r_2}{r_1}} & \text{if } r_1 > 0 \text{ and } r_2 < 0, \\ \sqrt{\frac{1}{r_1+r_2}} & \text{if } r_1 > 0 \text{ and } r_2 > 0, \end{cases} \quad (24.29)$$

$$\hat{q}_2 := \begin{cases} 1/\hat{q}_1 & \text{if } \frac{r_1}{r_2} < 0, \\ \hat{x}_1 & \text{if } \frac{r_1}{r_2} > 0 \end{cases}$$

and

$$r_1 := \frac{b_1(x_2 - x_3)}{b_3(x_2 - x_1)}, \quad r_2 := \frac{b_2(x_3 - x_1)}{b_3(x_2 - x_1)}.$$

Remark. If the points are ordered, i.e., $x_1 < x_2 < x_3$ (or alternatively in descending order), then either $A\hat{a} = (\hat{q}, 1/\hat{q}, \hat{q})^T$ or $A\hat{a} = (1/\hat{q}, \hat{q}, 1/\hat{q})^T$. This means that $\hat{\lambda}$ in Theorem 24.5.34 has alternating signs. In other words, the points $f(x_1)$, $f(x_3)$ lie above b_1 , b_3 and $f(x_2)$ lies below b_2 or conversely.

Later we will show that the alternating sign condition is true for general best polynomial approximation with respect to the Q -paranorm.

Corollary 24.5.37 is the basis of the Algorithm 24.6, which finds the optimal line with respect to three points in each step and chooses the next three points if the minimum corresponding to their line becomes larger.

Proposition 24.5.3 *The algorithm computes the line $f(x) = \hat{\beta} + \hat{\alpha}x$ which minimizes (24.28).*

Algorithm. (Best line approximation with respect to l_q)

Input: (x_i, b_i) , $i = 1, \dots, m$ of pairwise distinct points $x_i \in \mathbb{R}$ and $b_i > 0$

Set $i_1 := 1$, $i_2 := 2$ and $stop\text{signal} := -1$.

While $stop\text{signal} = -1$ do

1. For $i = 1, \dots, m$; $i \neq i_1, i_2$ compute

$$r_{1,i} := \frac{b_{i_1}(x_{i_2} - x_i)}{b_i(x_{i_2} - x_{i_1})}, \quad r_{2,i} := \frac{b_{i_2}(x_i - x_{i_1})}{b_i(x_{i_2} - x_{i_1})}.$$

2. Compute $\hat{a}_j = \max_i \{ \|\hat{x}_1(r_{1,i}, r_{2,i})\|_Q \}$ by (24.29). Let $j \neq i_1, i_2$ be an index where the maximum is attained and $\hat{x}_1 = \hat{x}_1(r_{1,j}, r_{2,j})$.

3. Compute $a := \max_i \{ \|r_{1,i}\hat{x}_1 + r_{2,i}\hat{x}_2\|_Q \}$.

Let k be an index, where the maximum is attained.

4. If $a \leq \hat{a}_j$ then $stop\text{signal} = 1$ and $\hat{a} = \hat{a}_j$,

$$\begin{pmatrix} \hat{\beta} \\ \hat{\alpha} \end{pmatrix} = \frac{1}{x_{i_2} - x_{i_1}} \begin{pmatrix} x_{i_2} & -x_{i_1} \\ -1 & 1 \end{pmatrix} \begin{pmatrix} b_{i_1} \hat{q}_1 \\ b_{i_2} / \hat{q}_1 \end{pmatrix},$$

otherwise set $i_1 := j$ and $i_2 := k$ and return to 1.

Figure 24.6: Algorithm finding best linear approximation under l_q .

Remark. Alternatively, one can deal with ordered points $b_1 < b_2 < b_3$ which restricts the effort in (24.29) to $\hat{q}_1 = \frac{r_2}{1-r_1}$ but requires an ascending ordering of the points x_{i_1}, x_{i_2}, x_j in each step of the algorithm.

Finally, we want to generalize the remark on the signs of the Lagrange multipliers given after Corollary 24.5.37. Remember that the set of polynomials $\Phi_i(x) = x^{i-1}$, $i = 1, \dots, n$ forms a Chebyshev set (see Def. 24.5.30). Applying again Lemma 24.5.33, one can easily prove the following result.

Theorem 24.5.38 *Let $\Phi_i : I \rightarrow \mathbb{R}$, $i = 1, \dots, n$ be a Chebyshev set and let $x_1 < \dots < x_{n+1}$ be points in I . Then, for*

$$\Phi := (\Phi_j(x_i))_{i,j=1}^{n+1,n},$$

the Lagrange multipliers $\hat{\lambda}_i$, $i = 1, \dots, n + 1$ corresponding to the minimizer of $\|\Phi \cdot\|_Q$ have alternating signs.

For our polynomial approximation problem $\operatorname{argmin}_{\vec{a} \in \mathbb{R}^n} \|A\vec{a}\|_Q$ with $A \in \mathbb{R}^{(n+1),n}$ defined by (24.21) and ordered points $x_1 < \dots < x_{n+1}$ we see that $A = \operatorname{diag}(1/b_i)_{i=1}^{n+1} \Phi$, where Φ is the matrix belonging to the Chebyshev set $\Phi_i(x) = x^{i-1}$. Since the b_i are positive, we obtain immediately that the Lagrange multipliers $\hat{\lambda}_i$ have alternating signs. Again, this means that $\hat{f}(x_i) - b_i$ has alternating signs.

Using Theorem 24.5.38, we not only simplify Algorithm 24.6 but also use Algorithm 24.5 even for approximation by arbitrary Chebyshev sets of size

EXC

$n > 2$. We only have to provide a routine which solves the following system of equations:

$$\begin{aligned} a_{1,1}x_1 + a_{1,2}x_2 + a_{1,n}x_n &= \lambda^{+1} \\ a_{2,1}x_2 + a_{2,2}x_2 + a_{2,n}x_n &= \lambda^{-1} \\ &\dots \quad \dots \\ a_{n+1,1}x_2 + a_{n+1,2}x_2 + a_{n+1,n}x_n &= \lambda^{(-1)^n} \end{aligned}$$

Let us illustrate this for $n = 2$. In this case, we can write

$$\begin{aligned} \frac{1}{\lambda}(\alpha + \beta x_1) &= y_1 \\ \lambda(\alpha + \beta x_2) &= y_2 \\ \frac{1}{\lambda}(\alpha + \beta x_3) &= y_3 \end{aligned}$$

If we number the above equations from 1 to 3, then we may conclude that

$$\begin{aligned} 3 &\implies & \alpha &= \lambda y_3 - \beta x_3 & (*) \\ 1, (*) &\implies & \lambda y_3 - \beta x_3 + \beta x_1 &= \lambda y_1 \\ &\implies & (y_3 - y_1)\lambda &= (x_3 - x_1)\beta \\ &\implies & \lambda &= \frac{x_3 - x_1}{y_3 - y_1}\beta & (**) \\ &\implies & \lambda &= q_{13}\beta & (**) \\ 2, (*), (**) &\implies & q_{13}\beta(q_{13}y_3\beta - \beta x_3 + \beta x_2) &= y_2 \\ &\implies & \beta^2(q_{13}y_3 - x_3 + x_2) &= y_2 q_{13}^{-1} \\ &\implies & \beta &= \sqrt{g^{-1}y_2 q_{13}^{-1}} \end{aligned}$$

where

$$\begin{aligned} q_{13} &:= \frac{x_3 - x_1}{y_3 - y_1} \\ g &:= q_{13}y_3 - x_3 + x_2 \end{aligned}$$

Caution is necessary, if $\beta = 0$. Then:

$$\begin{aligned} \beta &= 0 \\ \alpha &= \lambda y_1 \\ \lambda &= \sqrt{y_2/y_1} \end{aligned}$$

24.5.6 Non-Linear Models

In general there is a lot to say about this subject and we refer the reader to the literature [?]. However, let us consider two simple problems, which we can solve using algorithms we already know:

1. Find the best approximation using $e^{p(x)}$ under l_q , and
2. find the best approximation using $\ln(p(x))$ under l_∞ ,

where $p(x)$ is a linear combination of a set of Chebyshev functions.

Let us start with the first problem. That is, we ask for an exponential function

$$\hat{f} = e^{\sum_{j=1}^n \alpha_j \Phi_j}$$

which best approximates under l_q a given set of points (x_i, y_i) , $i = 1, \dots, m$ with pairwise distinct $x_i \in \mathbb{R}^d$ and $y_i > 0$, $1 \leq i \leq m$. Note that $\hat{f} > 0$ by definition. Since the \ln function increases strictly monotone this is equivalent to minimizing

$$\begin{aligned} \ln \left(\max_{i=1, \dots, m} \max \left\{ \frac{y_i}{\hat{f}(x_i)}, \frac{\hat{f}(x_i)}{y_i} \right\} \right) &= \max_{i=1, \dots, m} \max \{ \ln y_i - \ln \hat{f}(x_i), \ln \hat{f}(x_i) - \ln y_i \} \\ &= \max_{i=1, \dots, m} \left| \ln y_i - \sum_{j=1}^n \alpha_j \Phi_j(x_i) \right| \\ &= \|(\ln y_i)_{i=1}^m - \Phi \alpha\|_{\infty}. \end{aligned}$$

Thus, it remains to find the best function $\sum_{j=1}^n \alpha_j \Phi_j(x_i)$ with respect to the l_{∞} norm.

It is now easy to see that we can solve the second problem as follows. Let (x_i, y_i) be the data we want to approximate by a function of the form $\text{EXC} \ln(p(x))$ while minimizing the Chebyshev norm. We can do so by finding the best approximation of (x_i, e^{y_i}) under l_q .

Part V

Implementation

Chapter 25

Architecture of a Query Compiler

25.1 Compilation process

25.2 Architecture

Figure 25.1 a path of a query through the optimizer. For every step, a single component is responsible. Providing a facade for the components results in the overall architecture (Fig. 25.2). Every component is reentrant and stateless. The information necessary for a component to process a query is passed via references to control blocks. Control blocks are discussed next, then we discuss memory management. Subsequent sections describe the components in some detail.

25.3 Control Blocks

It is very convenient to have a hierarchy of control blocks within the optimizer. Figure 25.3 shows some of the control blocks. For simplification, those blocks concerned with session handling and transaction handling are omitted. Every routine call within the optimizer has a control block pointer as a parameter. The routines belonging to a specific phase have a pointer to the phase' specific control block as a parameter. For example, the routines in NFST have a `NFST_CB` pointer as a parameter. We now discuss the purpose of the different control blocks.

The global control block governs the behavior of the query compiler. It contains boolean variables indicating which phases to perform and which phases of the compilation process are to be traced. It also contains indicators for the individual phases. For example, for the first rewrite phase it contains indicators which rules to apply, which rules to trace and so on. These control indicators are manipulated by the driver which also allows to step through the different phases. This is very important for debugging purposes. Besides this overall control of the query compilers behavior, the global control block also contains

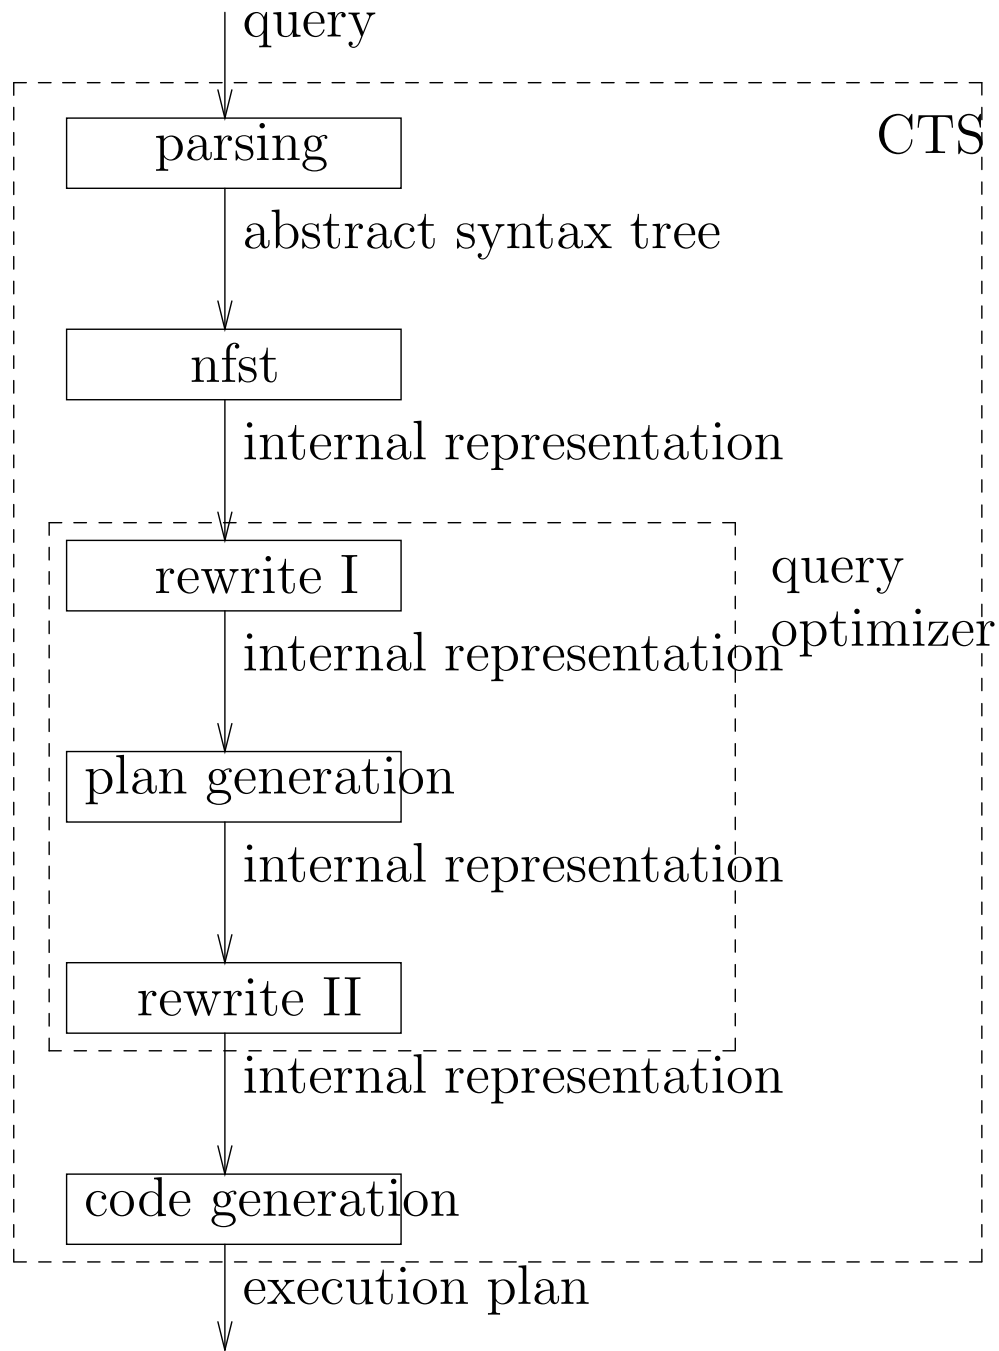


Figure 25.1: The compilation process

a pointer to the schema cache. The schema cache itself allows to look up type names, relations, extensions, indexes, and so on.

The query control block contains all the information gathered for the current query so far. It contains the abstract syntax tree, after its construction, the analyzed and translated query after NFST has been applied, the rewritten plan

after the `Rewrite.I` phase, and so on. It also contains a link to the memory manager that manages memory for this specific query. After the control block for a query is created, the memory manager is initialized. During the destructor call, the memory manager is destroyed and memory released.

Some components need helpers. These are also associated with the control blocks. We discuss them together with the components.

25.4 Memory Management

There are three approaches to memory management in query optimizers. The first approach is to use an automatic garbage collector if the language provides one. This is not necessarily the most efficient approach but by far the most convenient one. This approach can be imitated by an implementation based on smart pointers. I would not recommend doing so since the treatment of cycles can be quite tricky and it is inefficient. Another approach would be to collect all references to newly created objects and release these after the query has been processed. This approach is easy to implement, very convenient (transparent to the implementor), but inefficient. A better approach is to allocate bigger areas of memory by a memory manager. Factories¹ then use these memory chunks to generate objects as necessary. After the query has been processed, the chunks are freed.

Here, we consider only memory whose duration lasts for the processing of a single query. In general, we have more kinds of memory whose validity conforms to sessions and transactions.

25.5 Tracing and Plan Visualization

25.6 Driver

25.7 Bibliography

¹Design pattern.

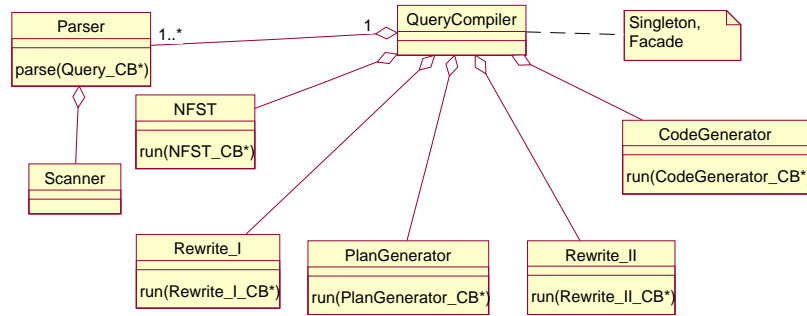


Figure 25.2: Class Architecture of the Query Compiler




Figure 25.3: Control Block Structure

Chapter 26

Internal Representations

26.1 Requirements

easy access to information

query representation: overall design goal: methods/functions with semantic meaning, not only syntactic meaning.

relationships: consumer/producer (occurrence) precedence order information equivalence of expressions (transitivity of equality) see also `expr.h` fuer andere funktionen/beziehungen die gebraucht werden

2-ebenen repraesentation. 2. ebene materialisiert einige beziehungen und funktionen, die haeufig gebraucht werden und kompliziert zu berechnen sind anderer grund fuer materialisierung: vermeide zuviele geschachtelte forschleifen. bsp: keycheck: gegeben eine menge von attributen und eine menge von schluesseln, ist die menge ein schluessel? teste jeden schluessel, kommt jedes element in schluessel in menge von attributen vor? (schon drei schleifen!!!)

modellierungsdetail: ein grosser struct mit dicken case oder feine klassenhierarchie. wann splitten: nur wenn innerhalb des optimierers verschiedene abarbeitung erfordert.

Representation: info captured: 1) 1st class information (information obvious in original query+(standard)semantic analysis) 2) 2nd class information (derived information) 3) historic information (during query optimization itself) - modified (original expression, modifier) - copied (original expression, copier) 4) information about the expression itself: (e.g.: `is_function_call`, `is_select`) 5) specific representations for specific purposes (optimization algorithms, code generation, semantic analysis) beziehungen zwischen diesen repraesentationen

info captured for 1) different parts of the optimizer

syntactic/semantic information

garbage collection: 1) manually 2) automatic 3) semi-automatic (collect references, free at end of query)

26.2 Algebraic Representations

relational algebra in: [167].

26.2.1 Graph Representations

26.2.2 Query Graph

also called *object graph*: [70, 840]

26.2.3 Operator Graph

used in: [714], [836]

enhanced to represent physical properties: [629]

with outerjoins: [626], [264]

graph representation and equivalence to calculus: [561]

26.3 Query Graph Model (QGM)

26.4 Classification of Predicates

klassifikation von praedikaten

- nach stelligkeit, wertigkeit (selektion, join, nasty)
- nach funktor (=, j, ..., between, oder boolsche funktion)
- nach funktion: fuer keys in index: start/stop/range/exact/enum range(in-predicate)
- nach sel-wert: simple (col = const), komplex (col = expr) cheap/expensive
- nach join wert: fuer hj, smj, hbnlj, ...
- korrelationspraedikate

26.5 Treatment of Distinct

26.6 Query Analysis and Materialization of Analysis Results

Questions:

1. was materialisieren wir
 - bsp: properties: zeiger auf property oder besser inline properties
 - bsp: unique number: entweder in expr oder getrennter dictionary struktur
2. was packen wir in die 1. repraesentation?
 - bsp: properties: zeiger auf property oder besser inline properties
 - bsp: unique number: entweder in expr oder getrennter dictionary struktur

```

query analysis (purpose, determine optimization algorithm)
#input relations, #predicates, #ex-quantifiers, #all-quantifiers,
#conjunctions, #disjunctions, #joininggraphkind(star,chain,tree,cyclic)
#strongly-connected-components (for crossproduct indication)
#false aggregates in projection list clause (implies grouping required)
/* remark: typical query optimizes should at least have two algorithms:
  - exhaustive (for large queries)
  - heuristic (for small queries)
*/

```

for blocks: indicator whether they should produce a null-tuple, in case they do not produce any tuple. this is nice for some rewrite rules. other possibility: if-statement in algebra.

26.7 Query and Plan Properties

Logical and Physical Properties of Plans

Ausführungsplänen können eine Reihe von Eigenschaften zugeordnet werden. Diese Eigenschaften fallen in drei Klassen

1. logische Eigenschaften, also beispielsweise
 - (a) beinhaltetete Relationen
 - (b) beinhaltetete Attribute
 - (c) angewendete Prädikate
2. physische Eigenschaften, also beispielsweise
 - (a) Ordnung der Tupel
 - (b) Strom oder Materialisierung des Ergebnisses
 - (c) Materialisierung im Hauptspeicher oder Hintergrundspeicher
 - (d) Zugriffspfade auf das Ergebnis
 - (e) Rechnerknoten des Ergebnis (im verteilten Fall)
 - (f) Kompression
3. quantitative Eigenschaften, also beispielsweise
 - (a) Anzahl der Elemente im Ergebnis
 - (b) Größe des Ergebnisses oder eines Ergebniselementes
 - (c) Auswertungskosten aufgeschlüsselt nach I/O, CPU und Kommunikationskosten

kosten: diese sind zu berechnen und dienen als grundlage fuer die planbewertung
 ges-kosten /* gesamt kosten (ressourcenverbrauch) */
 ges-kosten += cpu-instr / inst/sek
 ges-kosten += seek-kosten * overhead (waiting/cpu)
 ges-kosten += i/o-kosten * io-weight
 cpu-kosten /* reine cpu-kosten */
 i/o-kosten /* hintergrundspeicherzugriff (warten auf platte + cpu fuer seitenzugriffe) */

com-kosten /* kommunikation */ com-init /* initialisierungskosten fuer kommunikationsvorgang */ com-exit /* exitkosten fuer kommunikationsvorgang */ com-cptu /* kosten fuer jede transfereinheit (z.b. byte) waehrend eines kommunikationsvorgangs */

kostenstruktur koennte etwas sein, dass ges/cpu/io kosten enthaelt. ausserdem waeren kosten fuer rescanning interessant, falls dies notwendig ist (pufferprobleme, indexscan und dann faellt seite raus) weiteres interessantes kostenmass sind die kosten, bis das erste tupel berechnet wird.

dies sind die konstanten, die system-abhaengig sind. am besten sind, sie werden gemessen. Hardware: #cpu-instruktionen pro sekunde #cpu-instruktionen fuer block zugriff/transfer lesen/schreiben #cpu-instruktionen pro transfer init/send/exit init/receive/exit ms fuer seek/latency/transfer pro nK block

RTS-kosten #cpu-instruktionen fuer open/next/close fuer scan operatoren unter verschiedenen voraussetzungen:mit/ohne praedikat, mit/ohne projektion (entsprechend den avm programmen) #cpu-instruktionen fuer open/next/close fuer jeden alg operator, #cpu-instruktionen fuer funktionen/operationen/praedikate/avm-befehle

statistics: first/large physical page of a relation number of pages of a relation -i to estimate scan cost measured sequential scan cost (no interference/plenty interference)

-properties:

- menge der quns
- menge der attribute
- menge der praedikate
- ordnung
- boolean properties
- globale menge der gepipelineten quns
- kostenvektor
- cardinalitaeten bewiesen/geschaetzt
- gewuenschter puffer
- schluessel, fds
- #seiten, die durch ein fetch gelesen werden sollen
- menge der objekte, von denen der plan (der ja teilplan sein kann) abhaengt
- eigenschaften fuer parallele plaene
- eigenschaften fuer smp plaene

das folgende ist alles blabla. aber es weist auf den punkt hin,
das in dieser beziehung etwas getan werden muss.

```
--index: determine degree of clustering
- lese_rate = #gelesene_seiten / seiten_fuer_relation
  ein praedikate erniedrigt die lesen_rate, ein erneutes lesen aufgrund einer
  falls TIDs sortiert werden, muss fetch_ration erneut berechnet werden
- seiten koennen in gruppen z.b. auf einem zylinder zusammengefasst werden
  und mit einem prefetch befehl geholt werden. anzahl seeks abschaetzen
- cluster_ration(CR)
  CR = P(read(t) ohne page read) = (card - anzahl pagefetch)/card
    = (card - (#pagefetch - #page))/card

  das ist besonderer quark
- cluster_factor(CF)
  CF = P(avoid unnecessary pagefetch) = (pagefetch/maxpagefetch)
    = card - #fetch / card - #pageinrel

  das ist besonderer quark
index retrieval on full key => beide faktoren auf 100% setzen, da
  innerhalb eines index die TIDs pro key-eintrag sortiert werden.
```

Speicherung von Properties unter dynamischem Programmieren und Mem-
oization: Kosten und andere Eigenschaften, die nicht vom Plan abhängen,
können pro Planklasse gespeichert werden und brauchen nicht pro Plan gespe-
ichert zu werden.

26.8 Conversion to the Internal Representation

26.8.1 Preprocessing

26.8.2 Translation into the Internal Representation

26.9 Bibliography

Chapter 27

Details on the Phases of Query Compilation

27.1 Parsing

Lexical analysis is pretty much the same as for traditional compilers. However, it is convenient to treat keywords as soft. This allows for example for relation names like *order* which is a keyword in SQL. This might be very convenient for users since SQL has plenty (several hundreds) of keywords. For some keywords like *select* there is less danger of it being a relation name. A solution for *group* and *order* would be to lex them as a single token together with the following *by*.

Parsing again is very similar to parsing in compiler construction. For both, lexing and parsing, generators can be used to generate these components. The parser specification of SQL is quite lengthy while the one for OQL is pretty compact. In both cases, a LALR(2) grammar suffices. The outcome of the parser should be an abstract syntax tree. Again the data structure for abstract syntax trees (ast) as well as operations to deal with them (allocation, deletion, traversal) can be generated from an according ast specification.

During parsing already some of the basic rewriting techniques can be applied. For example, **between** can be eliminated.

In BD II, there are currently four parsers (for SQL, OQL, NQL (a clean version of XQuery), XQuery). The driver allows to step through the query compiler and allows to influence its overall behavior. For example, several trace levels can be switched on and off while within the driver. Single rewrites can be enabled and disabled. Further, the driver allows to switch to a different query language. This is quite convenient for debugging purposes. We used the Cocktail tools to generate the lexer, parser, ast, and NFST component.

27.2 Semantic Analysis, Normalization, Factorization, Constant Folding, and Translation

The NFST component performs (at least) four different tasks:




Figure 27.1: Expression

1. normalization of expressions,
2. factorization of common subexpressions,
3. semantic analysis, and
4. translation into the internal algebra-based query representation.

Although these are different tasks, a single pass over the abstract syntax tree suffices to perform all these tasks in one step.

Consider the following example query:

```

select e.name, (d.salary / d.budget) * 100
from Employee e, Department d
where e.salary > 100000 and e.dno = d.dno
  
```

The internal representation of the expression $(d.salary / d.budget) * 100$ in the query is shown in Fig. 27.1. It contains two operator nodes for the operations “*” and “/”. At the bottom, we find IU nodes. IU stands for Information Unit. A single IU corresponds to a variable that can be bound to a value. Sample IUs are attributes of a relation or, as we will see, intermediate results. In the query representation, there are three IUs. The first two IUs are bound to attribute values for the attributes *salary* and *budget*. The third IU is bound to the constant 100.

NFST routines can be implemented using a typical compiler generator tool. It is implemented in a rule-based language. Every rule matches a specific kind of AST nodes and performs an action. The ast tree is processed in post order.

The hierarchy for organizing different kinds of expressions is shown in Fig 27.2. Here is a list of useful functions:

- occurrence of expressions in another expression




Figure 27.2: Expression hierarchy

- for a given expression: compute the set of occurring (consumed, free) IUs
- for a given expression: compute the set of produced IUs
- for a given IU, retrieve the block producing the IU
- determine whether some block returns a single value only
- computation of the transitivity of predicates, especially equality to derive its equivalence classes.
- determine whether some expression produces a subset of another expression
- constant folding
- merge and/or (from e.g. binary to n-ary) and push not operations
- replace a certain expression by another one
- deep and shallow copy

These functions can be implemented either as member functions of expressions or according to visitor/collector/mutator patterns. For more complex functions (consumer/producer) we recommend the latter.

Some of these functions will be called quite frequently, e.g. the consumer/producer, precedence ordering, equivalence (transitivity of equality) functions. So it might be convenient to compute these relationships only once and then materialize them. Since some transformation in the rewrite phases are quite complex, a recomputation of these materialized functions should be possible since their direct maintenance might be too complex.

27.3 Normalization

Fig. 27.3 shows the result after normalization. The idea of normalization is to introduce intermediate IUs such that all operators take only IUs as arguments. This representation is quite useful.

27.4 Factorization

Common subexpressions are factorized by replacing them with references to some IU. For the expressions in TPCD query 1, the result is shown in Fig. 27.4.




Figure 27.3: Expression

Factorization is enabled by a factorization component that takes care of all expressions seen so far and the IUs representing these expressions. Every expression encountered by some NFST routine is passed to the factorization. The result is a reference to an IU. This IU can be a new IU in case of a new expression, or an existing IU in case of a common subexpression. The factorization component is available to the NFST routines via the NFST control block which is associated with a factorization component (Fig.25.3).

27.5 Constant Folding

27.6 Semantic analysis

The main purpose of semantic analysis is to attach a type to every expression. For simple expressions it is very similar to traditional semantic analysis in compiler construction. The only difference occurs for references to schema constructs. The schema is persistence and references to e.g. relations or named objects have to be looked up there. For performance reasons it is convenient to have a schema cache in order to cache frequently used references. Another aspect complicating semantic analysis a little is that collection types are frequently used in the database context. Their incorporation is rather straight forward but the different collection types should be handled with care.

As programming languages, query languages provide a block structure. Consider for example the SQL query

```

...
  select  a, b, c
  from    A, B
  where  d > e and f = g
...

```




Figure 27.4: Query 1

Consider the semantic analysis of d . Since SQL provides *implicit name look up*, we have to check (formerly analyzed) relations A and B whether they provide an attribute called d . If none of them provides an attribute d , then we must check the next upper SFW-block. If at least one of the relations A or B provides an attribute d , we just check that only one of them provides such an attribute. Otherwise, there would be an unallowed ambiguity. The blockwise look up is handled by block handler. For every newly encountered block (e.g. SFW block), a new block is opened. All identifiers analyzed within that block are pushed into the list of identifiers for that block. In case the query language allows for implicit name resolution, it might also be convenient to push all the attributes of an analyzed relation into the blocks list. The lookup is then performed blockwise. Within every block, we have to check for ambiguities. If the lookup fails, we have to proceed looking up the identifier in the schema. The handling of blocks and lookups is performed by the BlockHandler component attached to the control block of the NFST component (Fig. 25.3).

Another departure from standard semantic analysis are *false aggregates* as provided by SQL.

```
select avg(age)
from Students
```

I call *count(age)* a *false aggregate* since a true aggregate function operators on a collection of values and returns a single value. Here, the situation is different. The attribute *age* is of type integer. Hence, for the average function with signature $avg : \{int\} \longrightarrow int$ the semantic analysis would detect a typing error. The result is that we have to treat these false aggregates as special cases. This is (mostly) not necessary for query languages like OQL.

27.7 Translation

The translation step translates the original AST representation into an internal representation. There are as many internal query representations as there are query compiler. They all build on calculus expressions, operator graphs build over some algebra, or tableaux representations [766, 767]. A very powerful representation that also captures the subtleties of duplicate handling is the query graph model (QGM) [586].

The representation we use here is a mixture of a typed algebra and calculus. Algebraic expressions are simple operator trees with algebraic operators like selection, join, etc. as nodes. These operator trees must be correctly typed. For example, we are very picky about whether a selection operator returns a set or a bag. The expression that more resemble a calculus representation than an algebraic expression is the *SFWD block* used in the internal representation. We first clarify our notion of block within the query representation described here and then give an example of an SFWD block. A block is everything that produces variable bindings. For example a SFWD-block that pretty directly corresponds to a SFW-block in SQL or OQL. Other examples of blocks are quantifier expressions and grouping operators. A block has the following ingredients:

- a list of inputs of type collection of tuples¹ (labeled *from*)
- a set of expressions whose top is an IU (labeled *define*)
- a selection predicate of type bool (labeled *where*)

For *quantifier blocks* and *group blocks*, the list of inputs is restricted to length one. The SFWD-block and the grouping block additionally have a projection list (labeled *select*) that indicates which IUs are to be projected (i.e. passed to subsequent operators). Blocks are typed (algebraic) expressions and can thus be mixed with other expressions and algebraic operator trees.

An example of a SFWD-block is shown in Fig. 27.5 where dashed lines indicate the *produced-by* relationship. The graph corresponds to the internal

¹We use a quite general notion of tuple: a tuple is a set of variable (IU) bindings.

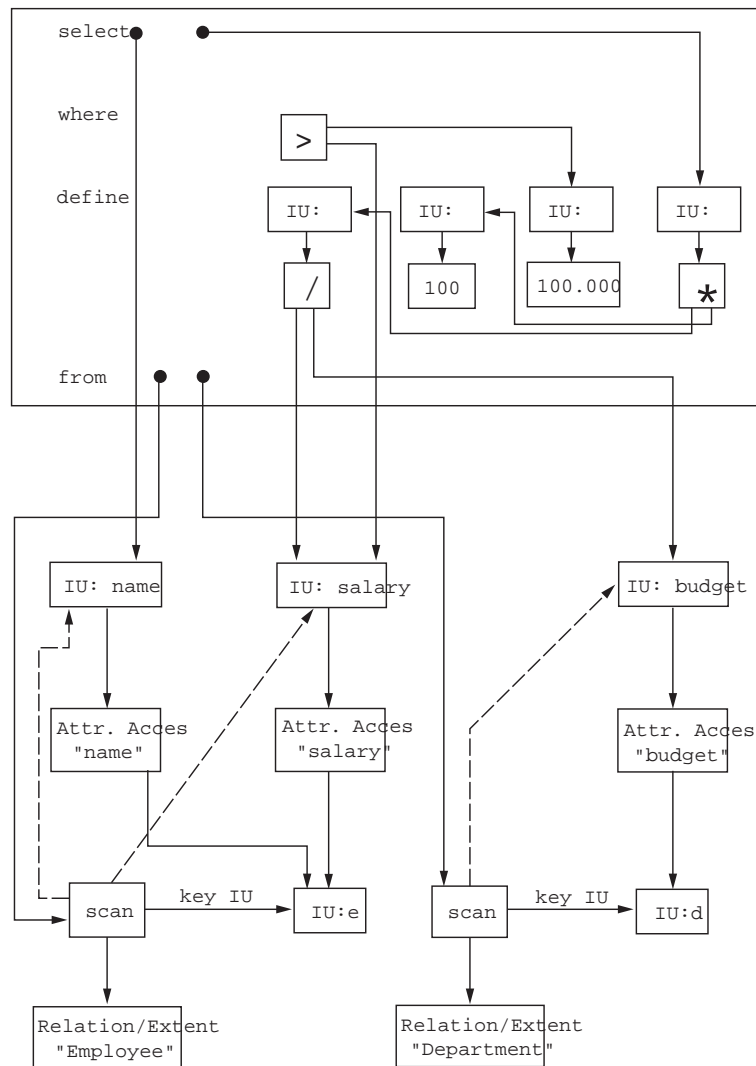


Figure 27.5: Internal representation

representation of our example query. The semantics of a SFWD-block can be described as follows. First, take the cross product of the collections of tuples found in the list of inputs. (If this is not possible, due to dependencies, d-joins have to be used.) Then, for every resulting tuple, compute the bindings for all the IUs mentioned in the *define* clause, apply the selection predicate and return all the bindings for the IUs mentioned in the *select* clause.

Although the SFWD-block looks neat, it lacks certain information that must be represented. This information concerns the role of the entries in the *from* clause and duplicate elimination. Let us start with the latter. There are three views relevant to duplicate processing:

1. the user view: did the user specify distinct?
2. the context view: does the occurrence or elimination of duplicates make

a difference for the query result?

3. the processing view: does the block produce duplicates?

All this information is attached to a block. This information can then be summarized to one of three values representing

- eliminate duplicates
- preserve duplicates
- don't care about duplicates
(The optimizer can feel free to do whatever is more efficient.)

This summary is also attached to every block. Let us illustrate this by a simple example:

```
select distinct ssno
from Employee
where ... and
        exists( select ... from ... where )
```

For the inner block, the user specifies that duplicates are to be preserved. However, duplicates or not does not modify the outcome of *exists*. Hence, the contextual information indicates that the outcome for the inner block is a don't care. The processing view can determine whether the block produces duplicates. If for all the entries in the *from* clause, a key is projected in the *select* clause, then the query does not produce duplicates. Hence, no special care has to be taken to remove duplicates produced by the outer block if we assume that *ssno* is the key of *Employee*.

No let us consider the annotations for the arguments in the *from* clause. The query

```
select distinct e.name
from Employee e, Department d
where e.dno = d.dno
```

retrieves only *Employee* attributes. Such a query is most efficiently evaluated by a semi-join. Hence, we can add a semi-join (SJ) annotation to the *Department d* clause.

For queries without a **distinct**, the result may be wrong (e.g. in case an employee works in several departments) since a typical semi-join just checks for existence. A special semi-join that preserves duplicates should be used. The according annotation is (SJ,PD). Another annotation occurs whenever an outer-join is used. Outer joins can (in SQL) be part of the *from* clause. Typically they have to be fully parenthesized since outer joins and regular joins not always commute. But under special circumstances, they commute and hence a list of entries in the *from* clause suffices [263]. Then, the entry to be preserved (the

outer part) should be annotated by (OJ). We use (AJ) as the anti-join annotation, and (DJ) for a d-join. To complete annotation, the case of a regular join can be annotated by (J). If the query language also supports all-quantifications, that translate to divisions, then the annotation (D) should be supported.

Since the graphical representation of a query is quite complex, we also use text representations of the result of the NFST phase. Consider the following OQL query:

```
select distinct s.name, s.age, s.supervisor.name, s.supervisor.age
from           s in Student
where         s.gpa > 8 and s.supervisor.age < 30
```

The annotated result (without duplicate annotations) of the normalization and factorization steps is

```
select distinct sn, sa, ssn, ssa
from           s in Student (J)
where         sg > 8 and ssa < 30
define       sn = s.name
              sg = s.gpa
              sa = s.age
              ss = s.supervisor
              ssn= ss.name
              ssa= ss.age
```

Semantic analysis just adds type information (which we never show).

In standard relational query processing multiple entries in the **from** clause are translated into a cross product. This is not always possible in object-oriented query processing. Consider the following query

```
select distinct s
from           s in Student, c in s.courses
where         c.name = "Database"
```

which after normalization yields

```
select distinct s
from           s in Student, c in s.courses
where         cn = "Database"
define       cn = c.name
```

The evaluation of *c in s.courses* is dependent on *s* and cannot be evaluated if no *s* is given. Hence, a cross product would not make much sense. To deal with this situation, the *d-join* has been introduced [161]. It is a binary operator that evaluates for every input tuple from its left input its right input and flattens the result. Consider the algebraic expression given in Fig. 27.6. For every student




Figure 27.6: An algebraic operator tree with a d-join

s from its left input, the d-join computes the set $s.courses$. For every course c in $s.courses$ an output tuple containing the original student s and a single course c is produced. If the evaluation of the right argument of the d-join is not dependent on the left argument, the d-join is equivalent with a cross product. The first optimization is to replace d-joins by cross products whenever possible.

Queries with a **group by** clause must be translated using the *unary grouping* operator GROUP which we denote by Γ . It is defined as

$$\Gamma_{g;\theta A;f}(e) = \{y.A \circ [g : G] | y \in e, \\ G = f(\{x | x \in e, x.A \theta y.A\})\}$$

where the subscripts have the following semantics: (i) g is a new attribute that will hold the elements of the group (ii) θA is the grouping criterion for a sequence of comparison operators θ and a sequence of attribute names A , and (iii) the function f will be applied to each group after it has been formed. We often use some abbreviations. If the comparison operator θ is equal to "=", we don't write it. If the function f is identity, we omit it. Hence, $\Gamma_{g;A}$ abbreviates $\Gamma_{g;=A;id}$.

Let us complete the discussion on internal query representation. We already mentioned algebraic operators like selection and join. These are called logical algebraic operators. Their implementations are called physical algebraic operators. Typically, there exist several possible implementations for a single logical algebraic operator. The most prominent example being the join operator with implementations like Grace join, sort-merge join, nested-loop join etc. All the operators can be modelled as objects. To do so, we extend the expression hierarchy by an algebra hierarchy. Although not shown in Fig 27.7, the algebra class should be a subclass of the expression class. This is not necessary for SQL but is a requirement for more orthogonal query languages like OQL.

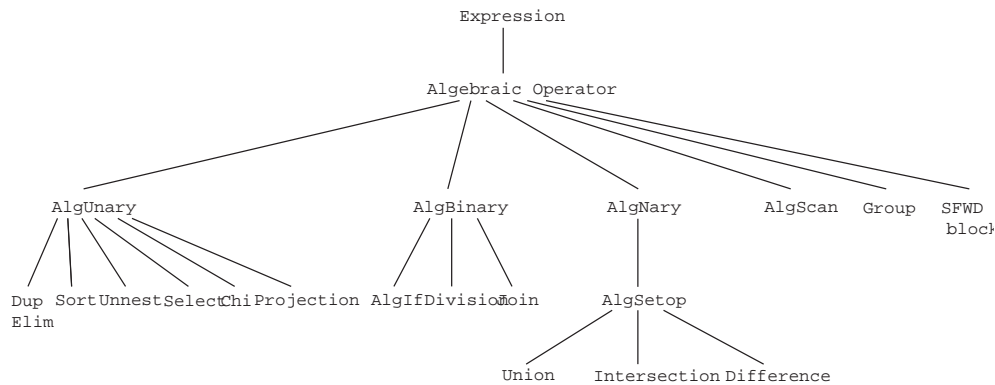


Figure 27.7: Algebra

27.8 Rewrite I

27.9 Plan Generation

27.10 Rewrite II

27.11 Code generation

In order to discuss the tasks of code generation, it is necessary to have a little understanding of the interface to the runtime system that interpretes the execution plan. I have chosen AODB as an example runtime system since this is one I know. The interface to AODB is defined by the AODB Virtual Machine (AVM). For simple operations, like arithmetic operations, comparisons and so on, AVM provides assembler-like operations that are interpreted at runtime. Simple AVM operations work on *registers*. A single register is able to hold the contents of exactly one IU. Additionally, AVM provides physical algebraic operators. These operators take AVM programs (possibly with algebraic operators) as arguments. There is one specialty about AVM programs though. In order to efficiently support factorization of common subexpressions involving arithmetic operations (as needed in aggregations like avg, sum), arithmetic operators in AVM can have two side effects. They are able to store the result of the operation into a register and they are able to add the result of the operation to the contents of another register. This is denoted by the result mode. If the result mode is A, they just add the result to some register, if it is C, they copy (store) the result to some register, if it is B, they do both. This is explored in the code for Query 1 of the TPC-D benchmark (Fig. 1.6).

Code generation has the following tasks. First it must map the physical operators in a plan to the operators of the AVM code. This mapping is a straight forward 1:1 mapping. Then, the code for the subscripts of the operators has to be generated. Subscripts are for example the predicate expressions for the selection and join operators. For grouping, several AVM programs have to be

generated. First program is the *init* program. It initializes the registers that will hold the results for the aggregate functions. For example, for an average operation, the register is initialized with 0. The *advance* program is executed once for every tuple to advance the aggregate computation. For example, for an average operations, the value of some register of the input tuple is added to the result register holding the average. The *finalize* program performs post-processing for aggregate functions. For example for the average, it divides the sum by the number of tuples. For hash-based grouping, the last two programs (see Fig.1.6) compute the hash value of the input register set and compare the group-by attributes of the input registers with those of every group in the hash bucket.

During the code generation for the subscripts factorization of common subexpression has to take place. Another task is register allocation and deallocation. This task is performed by the register manager. It uses subroutines to determine whether some registers are no longer needed. The register manager must also keep track in which register some IU is stored (if at all). Another component used during code generation is a factory that generates new AVM operations. This factory is associated with a table driven component that maps the operations used in the internal query representation to AVM opcodes.

27.12 Bibliography

Chapter 28

Hard-Wired Algorithms

28.1 Hard-wired Dynamic Programming

28.1.1 Introduction

Plan generation is performed block-wise. The goal is to generate a plan for every block. Typically, not all possible plans are generated. For example, the group operator (if necessary for the query) is mostly performed last (see also Sec. ??). This mainly leaves ordering joins and selections as the task of plan generation. A plan is an operator tree whose nodes consist of physical algebraic operators, e.g. selection, sort-operator, sort-merge and other joins, relation and index scans. The process of plan generation has received a lot of attention. Often, the term query optimization is used synonymously for the plan generation phase.

Figure 28.1 shows a plan for the block

```
select  e.name
from    Employee e, Department d
where   e.dno = d.dno and d.name = "shoe"
```

The bottom level contains two table scans that scan the base tables *Employee* and *Department*. Then, a selection operator is applied to restrict the departments to those named “shoe”. A nested-loop join is used to select those employees that work in the selected departments. The projection restricts the output to the name of the employees, as required by the query block. For such a plan, a *cost function* is used to estimate its cost. The goal of plan generation is to generate the cheapest possible plan. Costing is briefly sketched in Section ??.

The foundation of plan generation are algebraic equivalences. For e, e_1, e_2, \dots being algebraic expressions and p, q predicates, here are some example equiva-

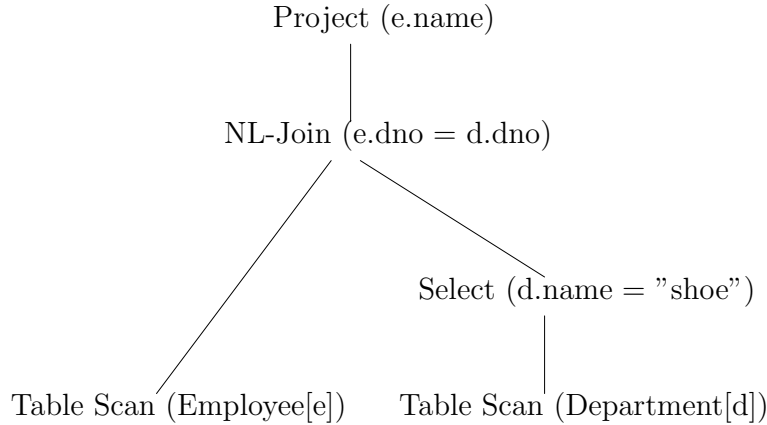


Figure 28.1: A sample execution plan

lences:

$$\begin{aligned}
 \sigma_p(\sigma_q(e)) &\equiv \sigma_q(\sigma_p(e)) \\
 \sigma_p(e_1 \bowtie_q e_2) &\equiv (\sigma_p(e_1)) \bowtie_q e_2 \quad \text{if } p \text{ is applicable to } e_1 \\
 e_1 \bowtie_p e_2 &\equiv e_2 \bowtie_p e_1 \\
 (e_1 \bowtie_p e_2) \bowtie_q e_3 &\equiv e_1 \bowtie_p (e_2 \bowtie_q e_3) \\
 e_1 \cup e_2 &\equiv e_2 \cup e_1 \\
 (e_1 \cup e_2) \cup e_3 &\equiv e_1 \cup (e_2 \cup e_3) \\
 e_1 \cap e_2 &\equiv e_2 \cap e_1 \\
 (e_1 \cap e_2) \cap e_3 &\equiv e_1 \cap (e_2 \cap e_3) \\
 \sigma_p(e_1 \cap e_2) &\equiv \sigma_p(e_1) \cap e_2
 \end{aligned}$$

For more equivalences and conditions that ought to be attached to the equivalences see the appendix 7.2. Note that commutativity and associativity of the join operator allow an arbitrary ordering. Since the join operator is the most expensive operation, ordering joins is the most prominent problem in plan generation.

These equivalences are of course independent of the actual implementation of the algebraic operators. The total number of plans equivalent to the original query block is called the *potential search space*. However, not always is the total search space considered. The set of plans equivalent to the original query considered by the plan generator is the *actual search space*. Since the System R plan generator [672], certain restrictions are applied. The most prominent are:

- Generate only plans where selections are pushed down as far as possible.
- Do not consider cross products if not absolutely necessary.
- Generate only left-deep trees.




Figure 28.2: Different join operator trees

- If the query block contains a grouping operation, the group operator is performed last.

Some comments are in order. Cross products are only necessary, if the query graph is unconnected where a query graph is defined as follows: the nodes are the relations and the edges correspond to the predicates (*boolean factors*¹) found in the **where** clause. More precisely, the query graph is a hypergraph, since a boolean factor may involve more than two relations. A left-deep tree is an operator tree where the right argument of a join operator always is a base relation. A plan with join operators whose both arguments are derived by other join operators is called *bushy tree*. Figure 28.2 gives an example of a left-deep tree and a bushy tree.

If we take all the above restrictions together, the problem boils down to ordering the join operators or relations. This problem has been studied extensively. The complexity of finding the best (according to some cost function) ordering (operator tree) was first studied by Ibaraki and Kameda [377]. They proved that the problem of generating optimal left-deep trees with no cross products is NP-hard for a special block-wise nested loop join cost function. This cost function applied in the proof is quite complex. Later it was shown that even if the cost function is very simple, the problem remains NP-hard [165]. The cost function (C_{out}) used there just adds up intermediate results sizes. This cost function is interesting in that it is the kernel of many other cost functions and it fulfills the ASI property of which we now have the following: If the cost function fulfills the *ASI property* and the query graph is acyclic, then the problem can be solved in polynomial time [377, 443]. Ono and Lohman gave examples that considering cross products can substantially improve performance [553]. However, generating optimal left-deep trees with cross products even for C_{out} makes the problem NP-hard [165]. Generating optimal bushy trees is

¹A *boolean factor* is a disjunction of basic predicates in a conjunctive normal form.

even harder. Even if there is no predicate, that is only cross products have to be used, the problem is NP-hard [657]. This is surprising since generating left-deep trees with cross products as the only operation is very simple: just sort the relations by increasing sizes.

Given the complexity of the problem, there are only two alternatives to generate plans: either explore the total search space or use heuristics. The former can be quite expensive. This is the reason why the above mentioned restrictions to the search space have traditionally been applied. The latter approach risks missing good plans. The best-known heuristics is to join the relation next, that results in the smallest next intermediate result. Estimating the cardinality of such results is discussed in Section ??.

Traditionally, selections were pushed as far down as possible. However, for expensive selection predicates (e.g. user defined predicates, those involving user-defined functions, predicates with subqueries) this does not suffice. For example, if a computer vision application has to compute the percentage of snow coverage for a given set of satellite images, this is not going to be cheap. In fact, it can be more expensive than a join operation. In these cases, pushing the expensive selection down misses good plans. That is why lately research started to take expensive predicates into account. However, some of the proposed solutions do not guarantee to find the optimal plans. Some approaches and their bugs are discussed in [132, 356, 354, 656, 658]. Although we will subsequently give an algorithm that incorporates correct predicate placement, not all plan generators do so. An alternative approach (though less good) is to pull-up expensive predicates in the Rewrite-II-phase.

There are several approaches to explore the search space. The original approach is to use dynamic programming [672]. The dynamic programming algorithm is typically hard-coded. Figure 28.3 illustrates the principle of bottom-up plan generation as applied in dynamic programming. The bottom level consists of the original relations to be joined. The next level consists of all plans that join a subset of cardinality two of the original relations. The next level contains all plans for subsets of cardinality three, and so on. With the advent of new query optimization techniques, new data models, extensible database systems, researchers were no longer satisfied with the hard-wired approach. Instead, they aimed for rule-based plan generation. There exist two different approaches for rule-based query optimizers. In the first approach, the algebraic equivalences that span the search space are used to transform some initial query plan derived from the query block into alternatives. As search strategies either exhaustive search is used or some stochastic approach such as simulated annealing, iterative improvement, genetic algorithms and the like [66, 382, 385, 386, 723, 747, 746, 749]. This is the *transformation-based* approach. This approach is quite inefficient. Another approach is to generate plans by rules in a bottom-up fashion. This is the *generation-based* approach. In this approach, either a dynamic programming algorithm [487] is used or memoization [306]. It is convenient to classify the rules used into logical and physical rules. The logical rules directly reflect the algebraic equivalences. The physical rules or implementation rules transform a logical algebraic operator into a physical algebraic operator. For example, a join-node becomes a nested-

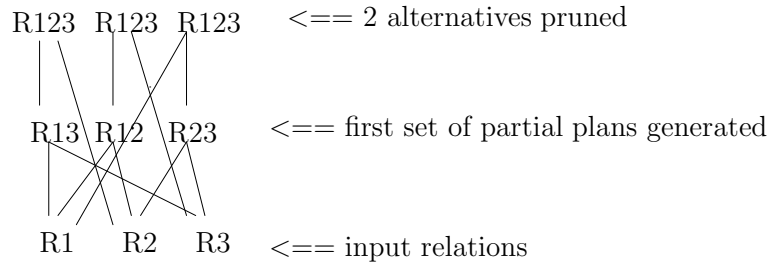


Figure 28.3: Bottom up plan generation

loop join node.

28.1.2 A plan generator for bushy trees

Within the brief discussion in the last subsection, we enumerated plans such that first all 1-relation plans are generated, then all 2-relation plans and so on. This enumeration order is not the most efficient one. Let us consider the simple problem where we have to generate exactly one best plan for the subsets of the n element set of relations to be joined. The empty subset is not meaningful, leaving the number of subsets to be investigated at $2^n - 1$. Enumerating these subsets can be done most efficient by enumerating them in *counting order*. That is, we initialize a n bit counter with 1 and count until have reached $2^n - 1$. The n bits represent the subsets. Note that with this enumeration order, plans are still generated bottom up. For a given subset R of the relations (encoded as the bit pattern a), we have to generate a plan from subsets of this subset (encoded as the bit pattern s). For example, if we only want to generate left-deep trees, then we must consider 1 element subsets and their complements. If we want to generate bushy trees, all subsets must be considered. We can generate these subsets by a very fast algorithm developed by Vance and Maier [778]:

```

s = a & -a;
while(s) {
    s = a & (s - a);
    process(s);
}

```

The meaning of *process(s)* depends on the kinds of plans we generate. If we concentrate on join ordering neglecting selection operations (i.e. pushing them) This step essentially looks up the plans for s and its complement \bar{s} and then joins the plans found there. Lookup is best implemented via an array with s as an index.

28.1.3 A plan generator for bushy trees and expensive selections

Figure 28.4 shows the pseudocode of a dynamic programming algorithm that generates plans with cross products, selections, and joins. It generates optimal

```

proc Optimal-Bushy-Tree( $R, P$ )
1  for  $k = 1$  to  $n$  do
2    for all  $k$ -subsets  $M_k$  of  $R$  do
3      for  $l = 0$  to  $\min(k, m)$  do
4        for all  $l$ -subsets  $P_l$  of  $M_k \cap R_S$  do
5           $best\_cost\_so\_far = \infty$ ;
6          for all subsets  $L$  of  $M_k$  with  $0 < |L| < k$  do
7             $L' = M_k \setminus L, V = P_l \cap L, V' = P_l \cap L'$ ;
8             $p = \bigwedge \{p_{i,j} \mid p_{i,j} \in P, R_i \in V, R_j \in V'\}$ ; //  $p=true$  might hold
9             $T = (T[L, V] \bowtie_p T[L', V'])$ ;
10           if  $Cost(T) < best\_cost\_so\_far$  then
11              $best\_cost\_so\_far = Cost(T)$ ;
12              $T[M_k, P_l] = T$ ;
13           fi;
14         od;
15         for all  $R \in P_l$  do
16            $T = \sigma_R(T[M_k, P_l \setminus \{R\}])$ ;
17           if  $Cost(T) < best\_cost\_so\_far$  then
18              $best\_cost\_so\_far = Cost(T)$ ;
19              $T[M_k, P_l] = T$ ;
20           fi;
21         od;
22       od;
23     od;
24   od;
25 od;
26 return  $T[R, S]$ ;

```

Figure 28.4: A Dynamic Programming Optimization Algorithm

bushy trees. Efficient implementation technique for the algorithm can be found in [778, 658]. As input parameters, the algorithm takes a set of relations R and a set of predicates P . The set of relations for which a selection predicate exists is denoted by R_S . We identify relations and predicates that apply to these relations. For all subsets M_k of the relations and subsets P_l of the predicates, an optimal plan is constructed and entered into the table T . The loops range over all M_k and P_l . Thereby, the set M_k is split into two disjoint subsets L and L' , and the set P_l is split into three parts (line 7). The first part (V) contains those predicates that apply to relations in L only. The second part (V') contains those predicates that apply to relations in L' only. The third part (p) is a conjunction of all the join predicates connecting relations in L and L' (line 8). Line 9 constructs a plan by joining the two plans found for the pairs $[L, V]$ and $[L', V']$ in the table T . If this plan has so far the best costs, it is memoized in the table (lines 10-12). Last, different possibilities of not pushing predicates in P_l are investigated (lines 15-19).

Another issue that complicates the application of dynamic programming are certain properties of plans. The most prominent such properties are *interesting orders* [672, 709, 710]. Take a look at the following query:

```
select    d.no, e.name
from      Employee e, Department d
where     e.dno = d.dno
order by d.dno
```

Here, the user requests the result to be order on *d.dno*. Incidentally, this is also a join attribute. During bottom up plan generation, we might think that a Grace hash join is more efficient than a sort-merge join since the cost of sorting the relations is too high. However, the result has to be sorted anyway so that this sort may pay off. Hence, we have to keep both plans. The approach is the following. In the example, an ordering on *d.dno* is called an interesting order. In general, any order that is helpful for ordering the output as requested by the user, for a join operator, for a grouping operator, or for duplicate elimination is called an *interesting order*. The dynamic programming algorithm is then modified such that plans are not pruned, if they produce different interesting orders.

28.1.4 A plan generator for bushy trees, expensive selections and functions

28.2 Bibliography

Chapter 29

Rule-Based Algorithms

29.1 Rule-based Dynamic Programming

The section is beyond the scope of the paper and the reader is referred to the starburst papers, especially [333, 462, 461, 487, 488].

29.2 Rule-based Memoization

This section is beyond the scope of the paper and the reader is referred to the Volcano and Cascade papers [292, 297, 302, 305, 306]. Both optimizer frameworks derived from the earlier Exodus query optimizer generator [290, 303].

29.3 Bibliography

Chapter 30

Example Query Compiler

30.1 Research Prototypes

30.1.1 AQUA and COLA

30.1.2 Black Dahlia II

30.1.3 Epoq

Für das objektorientierte Datenmodell *Encore* [844] wurde die Anfragesprache Equal [693, 692, 694], eine objektorientierte Algebra, die die Erzeugung von Objekten erlaubt, entwickelt. Zur Optimierung von Equal-Algebra-Ausdrücken soll der Optimierer Epoq dienen. Eine Realisierung von Epoq steht noch aus. Konkretisiert wurden jedoch bereits der Architekturansatz [525] und die Kontrolle der Alternativenerzeugung [523] innerhalb dieser Architektur. Einen Gesamtüberblick gibt die Dissertation von Mitchell [522].

Der Architekturvorschlag besteht aus einer generischen Architektur, die an einem Beispielloptimierer konkretisiert wurde [522, 523]. Die elementaren Bausteine der Architektur sind Regionen. Sie bestehen aus einer Kontrollkomponente und wiederum Regionen beziehungsweise Transformationen. Die einfachste Region ist dabei eine Transformation/Regel, die einen Algebraausdruck in einen äquivalenten Algebraausdruck umformt. Jede Region selbst wird wiederum als eine Transformation aufgefaßt. Innerhalb der Architektur werden nun diese Regionen in einer Hierarchie oder auch einem gerichteten azyklischen Graphen, organisiert. Abbildung 30.1 zeigt eine solche Beispielorganisation. Regionen selbst können bis auf die Kontrolle als Module im Sinne von Sciore und Sieg [668] aufgefaßt werden. Sie weisen sehr ähnliche Parameter und Schnittstellen auf. Während jedoch bei Sciore und Sieg die Kontrollstrategie eines Moduls aus einer festen Menge von gegebenen Kontrollstrategien ausgewählt werden muß, kann sie hier freier spezifiziert werden.

Unabhängig davon, ob die Transformationen einer Region wiederum Regionen sind oder elementare Transformationen, wird ihre Anwendung einheitlich von der Kontrolle der Region bestimmt. Die Aufgabe dieser Kontrolle besteht darin, eine Folge von Transformationen zu finden, die die gegebene Anfrage in eine äquivalente überführen. Sinngebend ist hierbei ein gewisses Ziel, das es zu erreichen gilt. Beispielsweise kann dieses Ziel lauten: Optimierte eine




Figure 30.1: Beispiel einer Epoq-Architektur

geschachtelte Anfrage. Um dieses Ziel zu erreichen, sind zwei grobe Schritte notwendig. Zunächst muß die Anfrage entschachtelt werden und als nächstes die entschachtelte Anfrage optimiert werden. Man sieht sofort, daß die Folge der Transformationen, die die Kontrolle auszuwählen hat, sowohl von den Eigenschaften der Anfrage selbst wie auch vom zu erfüllenden Ziel abhängt. Basierend auf dieser Beobachtung wird die Kontrolle nicht als Suchfunktion implementiert, sondern es wird das Planungsparadigma zur Realisierung gewählt. Die Kontrolle selbst wird mit Hilfe eines Satzes von Regeln spezifiziert, die aus Vorbedingung und Aktion bestehen.

Da es nicht möglich ist, im Vorfeld einen Plan, also eine Sequenz von Transformationen/Regionen, zu erstellen, der in garantierter Weise das Ziel erreicht, wird erlaubt, daß die Ausführung einer Transformation/Region fehlschlägt. In diesem Fall kann dann ein alternativer Plan erzeugt werden, der aber auf dem bisher Erreichten aufsetzt. Hierzu werden die Regeln, die die Kontrolle spezifizieren in Gruppen eingeteilt, wobei jeder Gruppe eine einheitliche Vorbedingung zugeordnet ist. Zu jeder Gruppe gehört dann eine Sequenz von Aktionen, die der Reihe nach ausprobiert werden. Schlägt eine vorangehende Aktion fehl, so wird die nächste in der Reihe der Aktionen angewendet. Schlagen alle Ak-

tionen fehl, so schlägt auch die Anwendung der Region fehl.

Jede Aktion selbst ist wiederum eine Sequenz von elementaren Aktionen. Jede dieser elementaren Aktionen ist entweder die Anwendung einer elementaren Transformation, der Aufruf einer Region oder der rekursive Aufruf des Planers mit einem neuformulierten Ziel, dessen Teilplan dann an entsprechender Stelle in die Aktion eingebaut wird.

Die Erweiterbarkeit dieses Ansatzes um neue Regionen scheint einfach möglich, da die Schnittstelle der Regionen genormt ist. Probleme könnte es lediglich bei den Kontrollstrategien geben, da nicht klar ist, ob die benutzte Regelsprache mächtig genug ist, um alle wünschenswerten Kontrollstrategien zu verwirklichen.

Die Frage, ob die einzelnen Komponenten des Optimierers, also die Regionen, evaluiert werden können, ist schwierig zu beantworten. Dafür spricht jedoch, daß jede Region in einem gewissen Kontext aufgerufen wird, also zur Erreichung eines bestimmten Zieles bei der Optimierung einer Anfrage mit ebenso bestimmten Eigenschaften. Beurteilen kann man daher die Erfolgsquote einer Region innerhalb ihrer verschiedenen Anwendungen. Da jede Region lediglich eine Alternative erzeugen darf, aufgrund des *eine Region ist eine Transformation*-Paradigmas, ist schwer zu sagen, in wie weit sich die durch die beschriebene Bewertung gewonnene Information zur Verbesserung der Regionen oder des Gesamtoptimierers einsetzen läßt.

Da auch hier der transformierende Ansatz zugrunde liegt, treffen die bereits diskutierten Probleme auch für den Optimierer für Straube zu.

Einen stetigen Leistungsabfall könnte man durch die Realisierung von alternativen Regionen erreichen, indem man ein Ziel *OptimiereSchnell* einführt, das dann entsprechend weniger sorgfältige, aber schnellere Regionen aufruft. Vorhersagen über der Güte (bei gegebener Optimierungszeit) scheinen aber schwerlich möglich.

30.1.4 Ereq

A primary goal of the EREQ project is to define a common architecture for the next generation of database managers. This architecture now includes

- * the query language OQL (a la ODMG), * the logical algebra AQUA (a la Brown), and * the physical algebra OPA (a la OGI/PSU).

It also includes

- * software to parse OQL into AQUA (a la Bolo)

and query optimizers:

- * OPT++ (Wisconsin), * EPOQ (Brown), * Cascades (PSU/OGI), and * Reflective Optimizer (OGI).

In order to test this architecture, we hope to conduct a "bakeoff" in which the four query optimizers will participate. The primary goal of this bakeoff is to determine whether optimizers written in different contexts can accommodate the architecture we have defined. Secondly, we hope to collect enough performance statistics to draw some conclusions about the four optimizers, which have been written using significantly different paradigms.

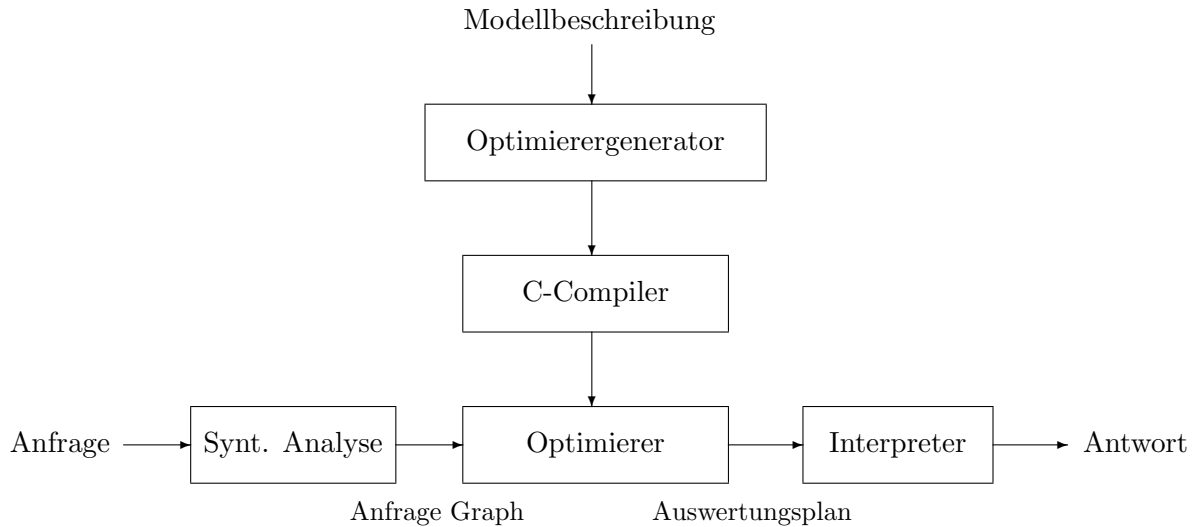


Figure 30.2: Exodus Optimierer Generator

At present, OGI and PSU are testing their optimizers on the bakeoff queries. Here is the prototype bakeoff optimizer developed at OGI. This set of Web pages is meant to report on the current progress of their effort, and to define the bakeoff rules. Please email your suggestions for improvement to Leo Fegaras fegaras@cse.ogi.edu. Leo will route comments to the appropriate author.

<http://www.cse.ogi.edu/DISC/projects/ereq/bakeoff/bakeoff.html>

30.1.5 Exodus/Volcano/Cascade

Im Rahmen des Exodus-Projektes wurde ein Optimierergenerator entwickelt [303]. Einen Überblick über den Exodus-Optimierergenerator gibt Abbildung 30.2. Ein Model description file enthält alle Angaben, die für einen Optimierer nötig sind. Da der Exodus-Optimierergenerator verschiedene Datenmodelle unterstützen soll, enthält dieses File zunächst einmal die Definition der verfügbaren *Operatoren* und *Methoden*. Dabei werden mit *Operatoren* die Operatoren der logischen Algebra bezeichnet und mit *Methoden* diejenigen der physischen Algebra, also die Implementierungen der Operatoren. Das Model description file enthält weiterhin zwei Klassen von Regeln. *Transformationen* basieren auf algebraischen Gleichungen und führen einen Operatorbaum in einen anderen über. *Implementierungsregeln* wählen für einen gegebenen Operator eine Methode aus. Beide Klassen von Regeln haben einen linken Teil, der mit einem Teil des aktuellen Operatorgraphen übereinstimmen muß, einen rechten Teil, der den Operatorgraphen nach Anwendung der Regel beschreibt, und eine Bedingung, die erfüllt sein muß, damit die Regel angewendet werden kann. Während die linke und rechte Seite der Regel als Muster angegeben werden, wird die Bedingung durch C-Code beschrieben. Auch für die Transformation lassen sich C-Routinen verwenden. In einer abschließenden Sektion des Model description files finden sich dann die benötigten C-Routinen.

Aus dem Model description file wird durch den Optimierergenerator ein C-Programm erzeugt, das anschließend übersetzt und gebunden wird. Das Ergebnis ist dann der Anfrageoptimierer, der in der herkömmlichen Art und Weise verwendet werden kann. Es wurde ein übersetzender Ansatz für die Regeln gewählt und kein interpretierender, da in einem von den Autoren vorher durchgeführten Experiment sich die Regelinterpretation als zu langsam erwiesen hat.

Die Regelarbeitung im generierten Optimierer verwaltet eine Liste OPEN, in der alle anwendbaren Regeln gehalten werden. Ein Auswahlmechanismus bestimmt dann die nächste anzuwendende Regel und entfernt sie aus OPEN. Nach deren Anwendung werden die hierdurch ermöglichten Regelanwendungen detektiert und in OPEN vermerkt. Zur Implementierung des Auswahlmechanismus werden sowohl die Kosten eines aktuellen Ausdrucks als auch eine Abschätzung des Potentials einer Regel in Betracht gezogen. Diese Abschätzung des Potentials berechnet sich aus dem Quotienten der Kosten für einen Operatorbaum vor und nach Regelanwendung für eine Reihe von vorher durchgeführten Regelanwendungen. Mit Hilfe dieser beiden Angaben, den Kosten des aktuellen Operatorgraphen, auf den die Regel angewendet werden soll, und ihres Potentials können dann Abschätzungen über die Kosten des erzeugten Operatorgraphen berechnet werden. Die Suchstrategie ist Hill climbing.

Der von den Autoren vermerkte Hauptnachteil ihres Optimierergenerators, den sie jedoch für alle transformierenden regelbasierten Optimierer geltend machen, ist die Unmöglichkeit der Abschätzung der absoluten Güte eines Operatorbaumes und des Potentials eines Operatorbaumes im Hinblick auf zukünftige Optimierungen. Dadurch kann niemals abgeschätzt werden, ob der optimale Operatorbaum bereits erreicht wurde. Erst nach Generierung aller Alternativen ist die Auswahl des optimalen Operatorbaumes möglich. Weiter bedauern es die Autoren, daß es nicht möglich ist, den A*-Algorithmus als Suchfunktion zu verwenden, da die Abschätzung des Potentials oder der Distanz zum optimalen Operatorgraphen nicht möglich ist.

Zumindest kritisch gegenüberstehen sollte man auch der Bewertung einzelner Regeln, da diese, basierend auf algebraischen Gleichungen, von zu feiner Granularität sind, als daß eine allgemeine Bewertung möglich wäre. Die erfolgreiche Verwendung des Vertauschens zweier Verbundoperationen in einer Anfrage bedeutet noch lange nicht, daß diese Vertauschung auch in der nächsten Anfrage die Kosten verringert. Die Hauptursache für die kritische Einstellung gegenüber dieser recht ansprechenden Idee ist, daß eine Regelanwendung zu wenig Information/Kontext berücksichtigt. Würde dieses Manko beseitigt, wären Regeln also von entschieden größerer Granularität, so erschiene dieser Ansatz vielversprechend. Ein Beispiel wäre eine Regel, die alle Verbundoperationen gemäß einer gegebenen Heuristik ordnet, also ein komplexer Algorithmus, der mehr Wissen in seine Entscheidungen einbezieht.

Graefe selbst führt einige weitere Nachteile des Exodus-Optimierergenerators an, die dann zur Entwicklung des Volcano-Optimierergenerators führten [305, 306]. Unzureichend unterstützt werden

- nicht-triviale Kostenmodelle,

- Eigenschaften,
- Heuristiken und
- Transformationen von Subskripten von algebraischen Operatoren in algebraische Operatoren.

Der letzte Punkt ist insbesondere im Bereich der Objektbanken wesentlich, um beispielsweise Pfadausdrücke in eine Folge von Verbundoperationen umwandeln zu können.

Im Volcano-Optimierergenerator werden algebraische Ausdrücke wieder in einen Operatorbaum umgewandelt. Wie im Exodus-Optimierergenerator wird der Optimierer wieder mit einer Menge von transformierenden und implementierenden Regeln beschrieben. Die Nachteile des transformierenden Ansatz werden somit geerbt. Eine Trennung in zwei Phasen, wie bei vielen Optimierern anzutreffen, ist für den Volcano-Optimierergenerator nicht notwendig. Der Entwickler des Optimierers hat die Freiheit, die Phasen selbst festzulegen. Die Probleme, die sonst bei der Kopplung der algebraischen mit der nicht-algebraischen Optimierung auftreten, können also vermieden werden. Die Behandlung der Eigenschaften erfolgt zielorientiert. Die in der Anfrage geforderten Eigenschaften (bspw. Sortierung), werden der Suchfunktion als Parameter übergeben, damit gezielt Pläne erstellt werden, die diese erfüllen. Wenn ein Operator oder eine Methode eingebaut wird, so wird darauf geachtet, daß diese noch nicht erfüllten Eigenschaften durch den Operator oder die Methode erzielt werden. Die geforderten Eigenschaften dienen wieder als Zielbeschreibung für die nachfolgenden Aufrufe der Suchfunktion. Zu diesen Eigenschaften gehören auch Kostengrenzen, mit denen die Suchfunktion dann einen Branch-and-bound-Algorithmus implementiert. Bevor ein Plan für einen algebraischen Ausdruck generiert wird, wird in einer Hash-Tabelle nachgeschaut, ob ein entsprechender Ausdruck mit den geforderten Eigenschaften bereits existiert. Dadurch wird Doppeltarbeit vermieden. Bei beiden Optimierergeneratoren werden die Forderungen nach stetigem Leistungsabfall, früher Bewertung von Alternativen und Evaluierbarkeit einzelner Komponenten nicht erfüllt.

30.1.6 Freytags regelbasierte System R-Emulation

[251] zeigt, wie man mit Hilfe eines regelbasierten Ansatzes den Optimierer von System R [672] emulieren kann. Die Eingabe besteht aus einem Lisp-ähnlichen Ausdruck:

```
(select <proj-list>
      <sel-pred-list>
      <join-pred-list>
      <table-list>)
```

Die Projektionsliste besteht aus Attributspezifikationen der Form

```
<rel-name>.<attr-name>
```




Figure 30.3: Organisation der Optimierung

Diese werden auch für die Selektionsprädikate und Joinprädikate verwendet. Die Algebra beinhaltet sowohl Operatoren der logischen als auch der physischen Algebra. In einzelnen gibt es Scan-, Sort-, Projektions, Verbundoperatoren in einer logischen und verschiedenen physischen Ausprägungen. Die Erzeugung der Auswertungspläne wird in verschiedene Schritte unterteilt, die wiederum in Teilschritte zerlegt sind (siehe Abb. 30.3). Zunächst erfolgt die Übersetzung in die logische Algebra. Hier werden Scan-Operatoren um die Relationen gebaut und Selektionen, die nur eine Relation betreffen, in die Scan-Operatoren eingebaut. Der zweite Schritt generiert Zugriffspläne, indem der Scan-Operator durch einen einfachen File-Scan (FSCAN) ersetzt wird, oder falls möglich, durch einen Index-Scan (ISCAN). Der dritte Schritt generiert zunächst verschiedene Verbund-Reihenfolgen und bestimmt anschließend die Verbund-Methoden. Sie in System R wird zwischen Sort-merge- und Nested-loop-join unterschieden.

Es werden keinerlei Aussagen über die Auswahl einer Suchstrategie gemacht. Ziel ist es vielmehr, durch die Modellierung des System R Optimierers mit Hilfe eines Regelsystems die prinzipielle Brauchbarkeit des regelbasierten Ansatzes nachzuweisen.

30.1.7 Genesis

Das globale Ziel des Genesisprojektes [52, 53, 54, 57] war es, die gesamte Datenbanksoftware zu modularisieren und eine erhöhte Wiederverwendbarkeit von Datenbankmodulen zu erreichen. Zwei Teilziele wurden hierbei angestrebt:

1. Standardisierung der Schnittstellen und
2. Formulierung der Algorithmen unabhängig von der DBMS-Implementierung.

Wir interessieren uns hier lediglich für die Erreichung der Ziele beim Bau von Optimierern [50, 55].

Die Standardisierung der Schnittstellen wird durch eine Verallgemeinerung von Anfragegraphen erreicht. Die Algorithmen selbst werden durch Transformationen auf Anfragegraphen beschrieben. Man beachte, daß dies nicht bedeutet, daß die Algorithmen auch durch Transformationsregeln implementiert werden. Regeln werden lediglich als Beschreibungsmittel benutzt, um die Natur der Wiederverwendbarkeit von Optimierungsalgorithmen zu verstehen.

Die Optimierung wird in zwei Phasen eingeteilt, die Reduktionsphase und die Verbundphase. Die Reduktionsphase bildet Anfragegraphen, die auf nicht reduzierten Datenmengen arbeiten, auf solche ab, die auf reduzierten Datenmengen arbeiten. Die Reduktionsphase orientiert sich also deutlich an den Heuristiken zum Durchschieben von Selektionen und Projektionen. Die zweite Phase bestimmt Verbundordnungen. Damit ist die in den Papieren beschriebene Ausprägung des Ansatzes sehr konservativ in dem Sinne, daß nur klassische Datenmodelle betrachtet werden. Eine Anwendung der Methodik auf objektorientierte oder deduktive Datenmodelle steht noch aus.

Folglich lassen sich nur die existierenden klassischen Optimierungsansätze mit diesen Mitteln hinreichend gut beschreiben. Ebenso lassen sich die existierenden klassischen Optimierer mit den vorgestellten Mitteln als Zusammensetzung der ebenfalls im Formalismus erfaßten Algorithmen beschreiben. Die Zusammensetzung selbst wird mit algebraischen Termersetzungen beschrieben. Durch neue Kompositionsregeln lassen sich dann auch neue Optimierer beschreiben, die andere Kombinationen von Algorithmen verwenden.

Durch die formale, implementierungsunabhängige Beschreibung sowohl der einzelnen Optimierungsalgorithmen als auch der Zusammensetzung eines Optimierers wird die Wiederverwendbarkeit von bestehenden Algorithmen optimal unterstützt. Wichtig dabei ist auch die Verwendung der standardisierten Anfragegraphen. Dieser Punkt wird allerdings aufgeweicht, da auch vorgesehen ist, verschiedene Darstellungen von Anfragegraphen zu verwenden [53]. Hierdurch wird die Wiederverwendung von Implementierungen von Optimierungsalgorithmen natürlich in Frage gestellt, da diese üblicherweise nur auf einer bestimmten Darstellung der Anfragegraphen arbeiten.

Wenn neue Optimierungsansätze entwickelt werden, so lassen sie sich ebenfalls im vorgestellten Formalismus beschreiben. Gleiches gilt auch für neue Indexstrukturen, da auch diese formal beschrieben werden [51, 56]. Nicht abzusehen ist, in wieweit der standardisierte Anfragegraph Erweiterungen standhält. Dies ist jedoch kein spezifisches Problem des Genesisansatzes, sondern gilt für alle Optimierer. Es ist noch offen, ob es gelingt, die Optimierungsalgorithmen so zu spezifizieren und zu implementieren, daß sie unabhängig von der konkreten Darstellung oder Implementierung der Anfragegraphen arbeiten. Der objektorientierte Ansatz kann hier nützlich sein. Es erhebt sich jedoch die Frage, ob bei Einführung eines neuen Operators die bestehenden Algorithmen so implemen-

tierbar sind, daß sie diesen ignorieren können und trotzdem sinnvolle Arbeit leisten.

Die Beschränkung auf zwei Optimierungsphasen, die Reduktions- und die Verbundphase, ist keine Einschränkung, da auch sie mittels Termersetzungsregeln festgelegt wurde, und somit leicht geändert werden kann.

Da die Beschreibungen des Optimierers und der einzelnen Algorithmen unabhängig von der tatsächlichen Implementierung sind, sind auch die globale Kontrolle des Optimierers und die lokalen Kontrollen der einzelnen Algorithmen voneinander losgelöst. Dieses ist eine wichtige Forderung, um Erweiterbarkeit zu erreichen. Sie wird oft bei regelbasierten Optimierern verletzt und schränkt somit deren Erweiterbarkeit ein.

Die Evaluierbarkeit, die Vorhersagbarkeit und die frühe Bewertung von Alternativen sind mit dem vorgestellten Ansatz nicht möglich, da die einzelnen Algorithmen als Transformationen auf dem Anfragegraphen aufgefaßt werden. Dieser Nachteil gilt jedoch nicht allein für den hier vorgestellten Genesisansatz, sondern generell für alle bis auf einen Optimierer. Es ist allerdings nicht absehbar, ob dieser Nachteil aus dem verwendeten Formalismus resultiert oder lediglich aus deren Konkretisierung bei der Modellierung bestehender Optimierer. Es ist durchaus möglich, daß der Formalismus mit leichten Erweiterungen auch andere Ansätze, insbesondere den generierenden, beschreiben kann.

Insgesamt handelt es sich beim Genesisansatz um einen sehr brauchbaren Ansatz. Leider hat er, im Gegensatz zur Regelbasierung, nicht genug Widerhall gefunden hat. Er hat höchst wahrscheinlich mehr Möglichkeiten, die Anforderungen zu erfüllen, als bisher ausgelotet wurde.

30.1.8 GOMbgo

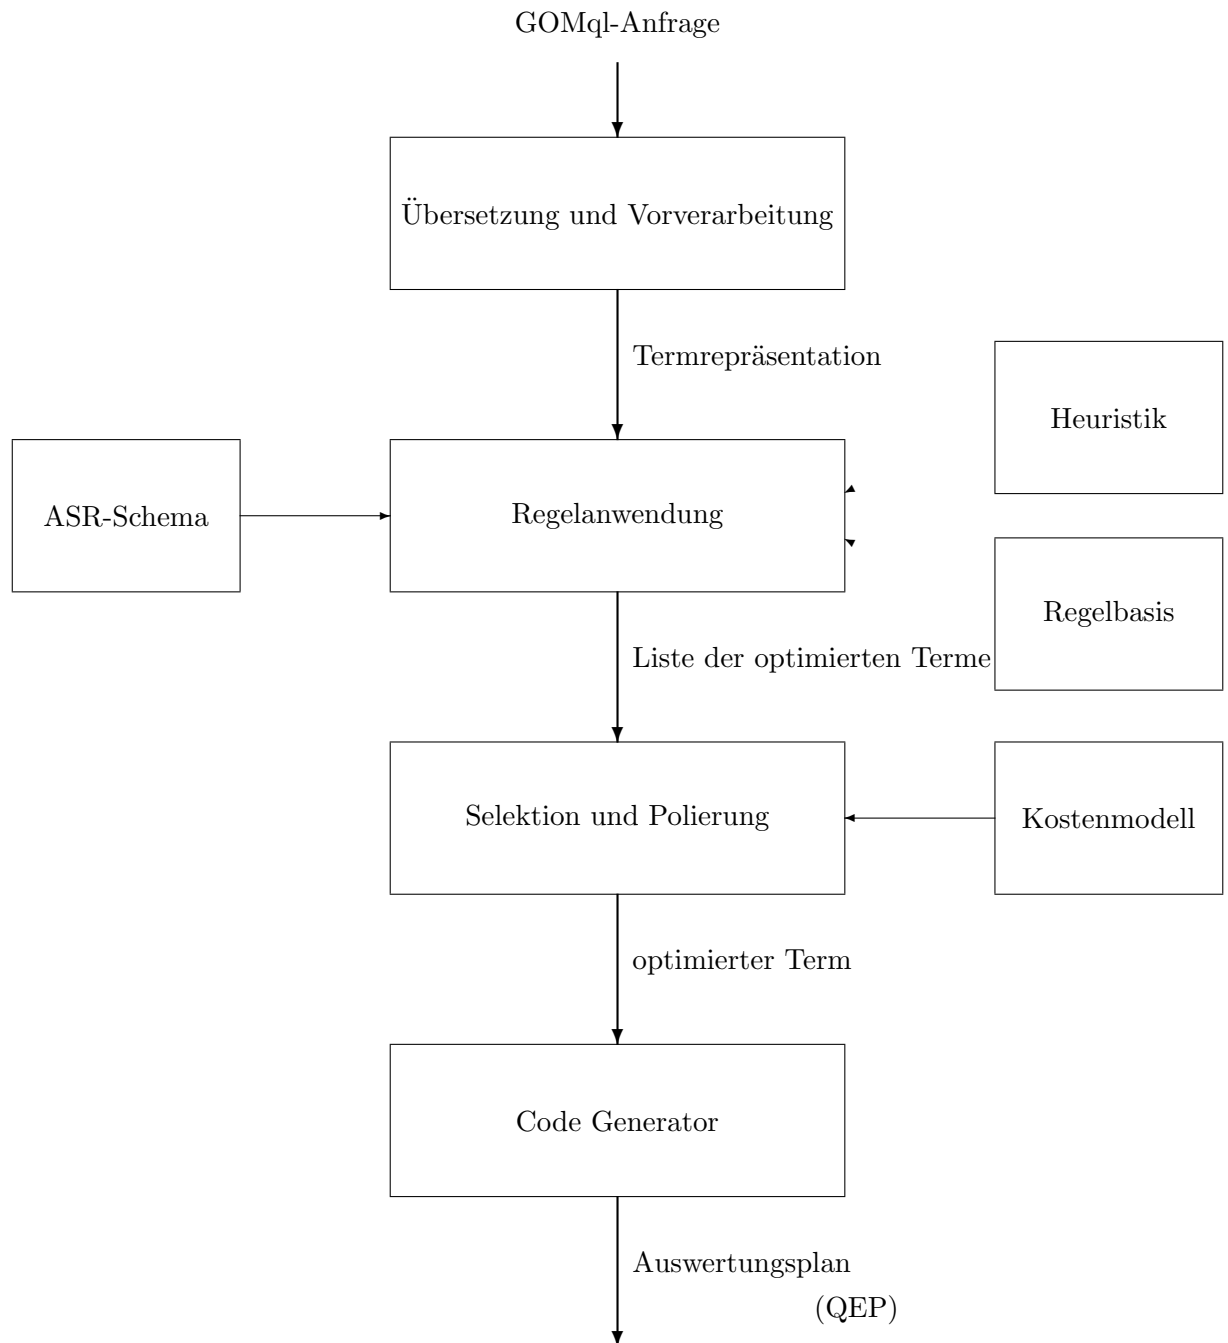


Figure 30.4: Ablauf der Optimierung

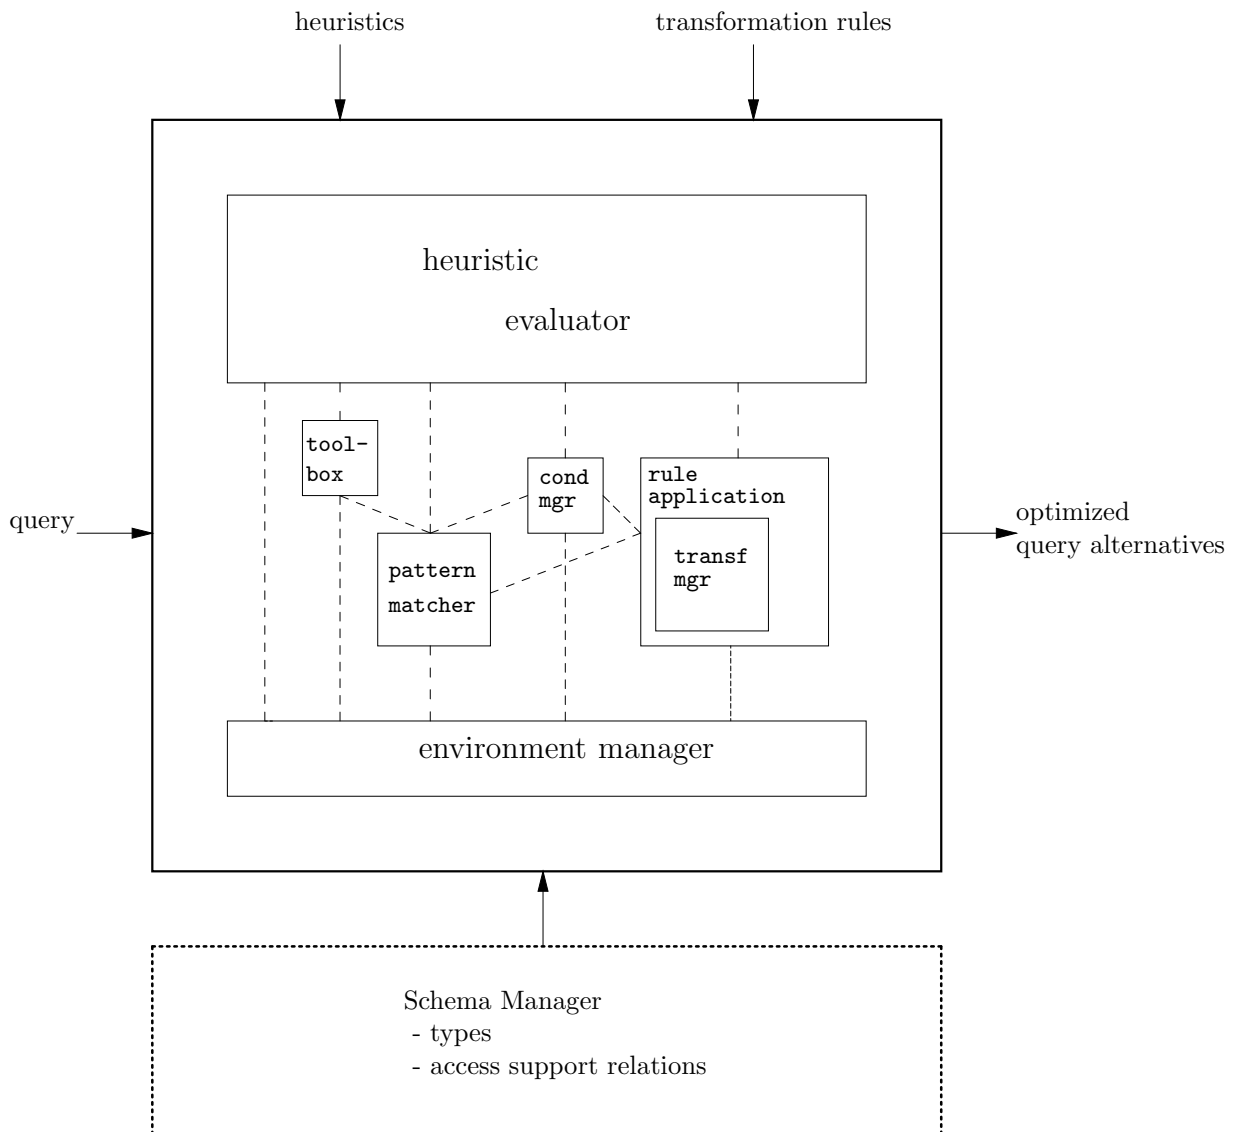


Figure 30.5: Architektur von GOMrbo

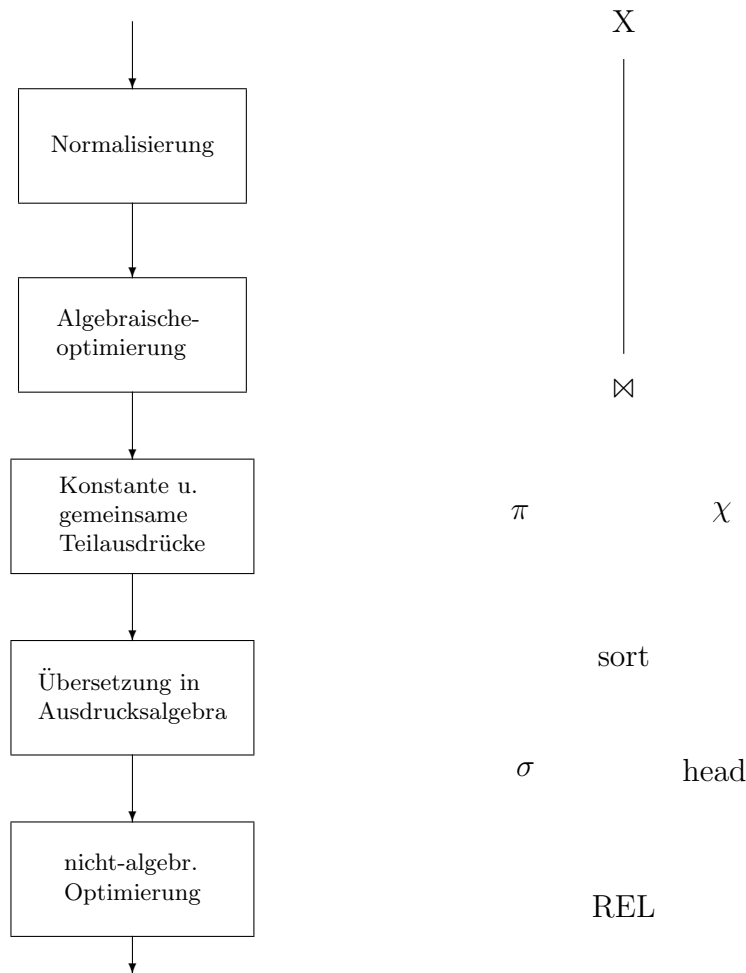


Figure 30.6: a) Architektur des Gral-Optimierers; b) Operatorhierarchie nach Kosten

30.1.9 Gral

Gral ist ein erweiterbares geometrisches Datenbanksystem. Der für dieses System entwickelte Optimierer, ein regelbasierter Optimierer in Reinkultur, erzeugt aus einer gegebenen Anfrage in fünf Schritten einen Ausführungsplan (s. Abb. 30.6 a) [60]. Die Anfragesprache ist gleich der verwendeten deskriptiven Algebra (*descriptive algebra*). Diese ist eine um geometrische Operatoren erweiterte relationale Algebra. Als zusätzliche Erweiterung enthält sie die Möglichkeit, Ausdrücke an Variablen zu binden. Ein Auswertungsplan wird durch einen Ausdruck der Ausführungsalgebra (*executable algebra*) dargestellt. Die Ausführungsalgebra beinhaltet im wesentlichen verschiedene Implementierungen der deskriptiven Algebra und Scan-Operationen. Die Trennung zwischen deskriptiver Algebra und Ausführungsalgebra ist strikt, das heißt, es kommen keine gemischten Ausdrücke vor (außer während der expliziten Konvertierung (Schritt 4)).

Die Schritte 1 und 3 sind durch feste Algorithmen implementiert. Während

der Normalisierung (Schritt 1) werden Variablenvorkommen durch die an sie gebundenen Ausdrücke ersetzt. Dies ist notwendig, um das Optimierungspotential vollständig erschließen zu können. Schritt 3 führt für konstante Ausdrücke Variablen ein. Die entspricht der Entschachtelung von Anfragen vom Typ N und A (s. Kapitel ?? und [424]). Die Behandlung von gemeinsamen Teilausdrücken ist noch nicht implementiert, aber für Schritt 3 vorgesehen.

Die Schritte 2, 4 und 5 sind regelbasiert. Zur Formulierung der Regeln wird eine Regelbeschreibungssprache (*rule description language*) verwendet. Die Beschreibungen der Regeln werden in einer Datei abgelegt. Innerhalb der Datei werden Regeln zu Gruppen (*sections*) zusammengefaßt. Diese Gruppen werden nacheinander angewandt. Daraus ergeben sich auch für einen Schritt mehrere kleinere Schritte. Beispielsweise ist der Schritt 2 im OPTEX-Optimierer für Gral in vier Teilschritte unterteilt:

1. Dekomposition von Selektionen mit komplexen Selektionsprädikaten in eine Folge von Selektionen mit einfachen Selektionsprädikaten und Zerlegung von Verbundoperationen in eine Folge von Selektionen und Kreuzprodukten.
2. Eigentlicher IMPROVING Schritt (siehe unten).
3. Teilausdrücke bestehend aus einer Selektion und einem unmittelbar folgenden Kreuzprodukt werden in Verbundoperationen umgewandelt.
4. Bestimmung einer Ordnung zwischen den Verbundoperationen und Kreuzprodukten. Dabei werden Kreuzprodukte zum Schluß ausgeführt und kleine Relationen zuerst verbunden.

Jeder Gruppe wird eine von drei in Gral implementierten Suchstrategien zugeordnet.

STANDARD Führt solange alle Regeln einer Gruppe aus, bis keine Regel mehr anwendbar ist. Es werden keine Vorkehrungen getroffen, um Endlosschleifen zu verhindern. Die Regeln müssen also dementsprechend formuliert werden. Diese Strategie kann für Schritte 2 und 5 verwendet werden.

IMPROVING Diese Strategie unterstützt algebraische Optimierung in der deskriptiven Algebra (Schritt 2). Das Ziel ist hierbei eine gute Ordnung der algebraischen Operatoren zu erlangen. Hierzu wird eine partielle Ordnung der algebraischen Operatoren gemäß ihrer Kosten definiert (s. Abb. 30.6 b) für ein Beispiel). Die IMPROVING-Strategie versucht dann die hierdurch definierte Ordnung in einem gegebenen Ausdruck zu erreichen. Hierzu wird sie zunächst rekursiv auf alle Teilausdrücke eines Ausdrucks angewendet. Regeln zur Umformung werden dann angewendet, wenn dadurch eine höhere Kohärenz der Operatorfolge im Ausdruck mit der der Operatorkostenhierarchie erreicht werden kann. Dies entspricht einem Bubble-sort auf dem Ausdruck. Ausdrücke mit der kleinsten Anzahl von *runs* werden bevorzugt. Dabei ist ein *run* eine Folge von Operatoren innerhalb des zu optimierenden Ausdrucks, dessen Operatoren gemäß der Operatorkostenhierarchie geordnet sind.

TRANSLATION Regelgruppen mit dieser Strategie werden während der Übersetzung von der deskriptiven Algebra in die Ausführungsalgebra angewendet (Schritt 4). Jede Regel beschreibt dabei die Übersetzung eines einzelnen deskriptiven Operators in einen Ausdruck der Ausführungsalgebra, also einen, der keine deskriptiven Operatoren enthalten darf. Die Übersetzung erfolgt lokal. Für Parameter, also beispielsweise Selektions- und Verbundprädikate, können Regeln angegeben werden, die einen Suchraum für die Reorganisation des Parameters erlauben. Hiermit kann man beispielsweise alle Permutationen einer Konjunktion erzeugen. Die Suchstrategie für die Parameterbestimmung ist erschöpfend und trägt Vorsorge, daß keine Zyklen auftreten. Eine Auswahl kann mittels des *valuation*-Eintrags in den Regeln getroffen werden. Dieser kann beispielsweise Kosten repräsentieren. Dementsprechend werden dann Regeln mit der kleinsten *valuation* bevorzugt. Jede für einen Parameter generierte Darstellung wird übersetzt.

Die Syntax für eine Regel ist

```

specification
definition
RULE
    pattern
    → result1 valuation1 if condition1
    ...
    → resultn valuationn if conditionn

```

wobei

specification von der Form

$$\text{SPEC spec}_1, \dots, \text{spec}_n$$

ist. Dabei sind die spec_i Range-Spezifikationen wie beispielsweise $op_i \text{ in } \langle OpSet \rangle$.

definition Variablen definiert (bspw. für Attributsequenzen). In Gral existieren verschiedene Sorten von Variablen für Attribute, Operationen, Relationen etc.

pattern ein Muster in Form eines Ausdrucks ist, der Variablen und Konstanten enthalten kann. Der Ausdruck kann ein Ausdruck der deskriptiven Algebra oder der Ausführungsalgebra sein.

condition_i eine Bedingung ist. Diese Bedingung ist ein allgemeiner boolescher Ausdruck. Spezielle Prädikate wie *ExistsIndex* (existiert ein Index für eine Relation?) werden von Gral zur Verfügung gestellt.

result_i wiederum ein Ausdruck ist, der das Ergebnis der Regel beschreibt.

valuation_i ist ein arithmetischer Ausdruck, der einen numerischen Wert zurückliefert. Dieser kann in einer (Gral unterstützt mehrere) Auswahlstrategie herangezogen werden: Es wird die Regel mit der kleinsten *valuation* bevorzugt.

Die Auswertung einer Regel erfolgt standardmäßig. Sei E der Ausdruck auf den die Regel angewendet werden soll.

```

if  $\exists$  Substitution  $\sigma$ , Unterausdruck  $E'$  von  $E$  mit  $E'\sigma = pattern$ 
  and  $\forall 1 \leq i \leq j: \neg condition_i$ 
  and  $condition_j$ 
then ersetze  $E'$  in  $E$  durch  $result_j\sigma$ 

```

Der Gral-Optimierer ist ein reiner regelbasierter Optimierer, der den Transformationsansatz verfolgt. Dementsprechend treffen alle vorher identifizierten Nachteile derselben zu.

Zu bemängeln sind im einzelnen folgende Punkte:

- Es erfolgt keine frühzeitige Bewertung der Alternativen.
- Die Suchstrategien sind fest eingebaut und nicht sonderlich ausgefeilt.
- Der Einbau von hochspezialisierten Algorithmen, die besondere Optimierungstechniken repräsentieren, ist schwierig, wenn nicht unmöglich.
- Eine Bestimmung der Verbundreihenfolge gemäß eines komplexeren Algorithmus ist nicht möglich.
- Da die Übersetzung in die Ausführungsalgebra lokal ist und keine Annotationen zugelassen sind, können vorhandene Sortierreihenfolgen nur schwer ausgenutzt werden.

Es wird nur eine Alternative der algebraischen Optimierung zur physischen Optimierung übergeben. Das kann zu Fällen führen, in denen der Optimierer niemals das Optimum finden kann. Wengleich dies auch im allgemeinen nicht immer möglich ist, so sollte jedoch diese Eigenschaft nicht inhärent sein.

Positiv zu vermerken ist, daß für IMPROVING und TRANSLATION der Aufwand für das Pattern-matching vermutlich gering gehalten werden kann.

30.1.10 Lambda-DB

<http://lambda.uta.edu/lambda-DB/manual/overview.html>

30.1.11 Lanzelotte in short

Query Language Der Lanzelotte-Optimierer verwendet keine spezielle Anfragesprache. Ausgangspunkt der Betrachtungen sind sog. Anfragegraphen (request graphs, query graphs). Einzelheiten stehen in meiner Ausarbeitung. In einem Papier ([451]) wird gezeigt wie man von einer Regelsprache (RDL) zu Anfragegraphen kommt.

Internal Representation Die interne Repraesentation einer Anfrage ist der Class Connection Graph. Dort enthalten sind die Datenbankobjekte (Extensionen), die in der Anfrage referenziert werden aus der Sicht des physikalischen Schemas und die in der Anfrage bedeutsamen Beziehungen zwischen diesen Extensionen (Joins, Attributpfade, Selektionen).

Query Execution Plans QEPs werden als (deep) processing trees repräsentiert.

Architecture Der Lanzelotte-Optimierer ist regelbasiert.

Transformation versus generation Lanzelotte bietet Regeln fuer beide Spielarten. Sie unterscheidet enumerative search (Generierung), randomized search (Transformation) und genetic search (Transformation).

Control/Search-Strategy Lanzelotte versucht von den Einzelheiten der verwendeten Strategien zu abstrahieren und stellt eine erweiterbare Optimierung vor, die die Einzelheiten ueberdeckt. Die tatsaechlich zu einem bestimmten Zeitpunkt verwendete Strategie wird durch "assertions" bestimmt. (Dazu steht nicht viel in den Papieren, vielleicht meint sie auch die Bedingungsteile der Regeln)

Cost Model Ziemlich aehnlich dem, das wir verwenden. Sie benutzt auch solche Sachen wie $card(C)$, $size(C)$, $ndist(A_i)$, $fan(A_i)$, $share(A_i)$. Einzelheiten stehen in meiner Ausarbeitung.

30.1.12 Opt++

wisconsin

30.1.13 Postgres

Postgres ist kein Objektbanksystem sondern fällt in die Klasse der erweiterten relationalen Systeme [733]. Die wesentlichen Erweiterungen sind

- berechenbare Attribute, die als Quel-Anfragen formuliert werden [731],
- Operationen [729],
- abstrakte Datentypen [728] und
- Regeln [732].

Diese beiden Punkte sollen uns jedoch an dieser Stelle nicht interessieren. Die dort entwickelten Optimierungstechniken, insbesondere die Materialisierung der berechenbaren Attribute, sind in der Literatur beschrieben [400, 344, 342, 343]. Unser Interesse richtet sich vielmehr auf eine neuere Publikation, in der eine Vorschlag für die Reihenfolgebestimmung von Selektionen und Verbundoperationen unterbreitet wird [356]. Diese soll im folgenden kurz vorgestellt werden. Zunächst jedoch einige Vorbemerkungen.

Wenn man eine Selektion verzögert, also nach einem Verbund ausführt, obwohl dies nicht notwendig wäre, so kann es passieren, daß das Selektionsprädikat auf mehr Tupeln ausgewertet werden muß. Es kann jedoch nicht passieren, daß es auf mehr verschiedenen Werten ausgeführt werden muß. Im Gegenteil, die Anzahl der Argumentewerte wird durch einen Verbund im allgemeinen verkleinert. Cached man also die bereits errechneten Werte des Selektionsprädikates, so wird die Anzahl der Auswertungen des Selektionsprädikates

nach einem Verbund zumindest nicht größer. Die Auswertung wird dann durch ein Nachschlagen ersetzt. Da wir hier nur teure Selektionsprädikate betrachten, ist ein Nachschlagen sehr billig gegenüber der Auswertung. Die Kosten für das Nachschlagen können sogar vernachlässigt werden. Es bleibt das Problem der Größe des Caches. Liegt Eingabe sortiert nach den Argumenten des Selektionsprädikates vor, so kann der die Größe des Caches unter Umständen auf 1 reduziert werden. Er erübrigt sich ganz, wenn man eine indirekte Repräsentation des Verbundergebnisses verwendet. Eine mögliche indirekte Repräsentation ist in Abbildung ?? dargestellt, wobei die linke der abgebildeten Relationen die Argumente für das betrachtete Selektionsprädikat enthalte.

Für jedes Selektionsprädikat $p(a_1, \dots, a_n)$ mit Argumenten a_i bezeichne c_p die Kosten der Auswertung auf einem Tupel. Diese setzen sich aus CPU- und I/O-Kosten zusammen (s. [356]). Ein *Plan* ist ein Baum, dessen Blätter *scan*-Knoten enthalten und dessen innere Knoten mit Selektions- und Verbundprädikaten markiert sind. Ein *Strom* in einem Plan ist ein Pfad von einem Blatt zur Wurzel. Die zentrale Idee ist nun die Selektions- und Verbundprädikate nicht zu unterscheiden, sondern gleich zu behandeln. Dabei wird angenommen, daß alle diese Prädikate auf dem Kreuzprodukt aller Relationen der betrachteten Anfrage arbeiten. Dies erfordert eine Anpassung der Kosten. Seien a_1, \dots, a_n die Relationen der betrachteten Anfrage und p ein Prädikat über den Relationen a_1, \dots, a_k . Dann sind die *globalen Kosten* von p wie folgt definiert:

$$C(p) = \frac{c_p}{\prod_{i=k+1}^n |a_i|}$$

Die globalen Kosten berechnen die Kosten der Auswertung des Prädikates über der gesamten Anfrage. Hierbei müssen natürlich diejenigen Relationen herausgenommen werden, die das Prädikat nicht beeinflussen. Zur Illustration nehme man an, p sei ein Selektionsprädikat auf nur einer Relation a_1 . Wendet man p direkt auf a_1 an, so entstehen die Kosten $c_p * |a_1|$. Im vereinheitlichten Modell wird angenommen, daß jedes Prädikat auf dem Kreuzprodukt aller in der Anfrage beteiligten Relationen ausgewertet wird. Es entstehen also die Kosten $C(p) * |a_1| * |a_2| * \dots * |a_n|$. Diese sind aber gleich $c_p * |a_1|$. Dies ist natürlich nur unter der Verwendung eines Caches für die Werte der Selektionsprädikate korrekt. Man beachte weiter, daß die Selektivität $s(p)$ eines Prädikates p unabhängig von der Lage innerhalb eines Stroms ist.

Der *globale Rang* eines Prädikates p ist definiert als

$$\text{rank}(p) = \frac{s(p)}{C(p)}$$

Man beachte, daß die Prädikate innerhalb eines Stroms nicht beliebig umordbar sind, da wir gewährleisten müssen, daß die von einem Prädikat benutzten Argumente auch vorhanden sein müssen. In [356] wird noch eine weitere Einschränkung vorgenommen: Die Verbundreihenfolge darf nicht angetastet werden. Es wird also vorausgesetzt, daß eine optimale Verbundreihenfolge bereits bestimmt wurde und nur noch die reinen Selektionsprädikate verschoben werden dürfen.

Betrachtet man zunächst einmal nur die Umordnung der Prädikate auf einem Strom, so erhält man bedingt durch die Umordbarkeitseinschränkungen

das *Sequentialisierungsproblem mit Vorrangbedingungen* für das Algorithmus mit Laufzeit $O(n \log n)$ (n ist die Stromlänge) eine optimale Lösung bekannt ist [533].

Das in [356] vorgeschlagene Verfahren wendet diesen Algorithmus solange auf jeden Strom an, bis keine Verbesserung mehr erzielt werden kann. Das Ergebnis ist ein polynomialer Algorithmus, der die optimale Lösung garantiert. Dies jedoch nur unter der Einschränkung, daß die Kosten des Joins linear sind.

Damit sind wir bereits bei einem der Nachteile des Verfahrens: Die Kosten der Verbundoperation nicht mitunter nicht linear sondern sogar quadratisch. Ein weiterer Nachteil liegt in der Voraussetzung, daß die optimale Verbundreihenfolge schon bestimmt wurde, denn diese hängt wesentlich davon ab, an welcher Stelle die Selektionen eingebaut werden. Üblicherweise wird bei der Bestimmung der optimalen Verbundreihenfolge vorausgesetzt, daß alle Selektionsprädikate soweit wie möglich nach unten verschoben werden. Dies ist jedoch jetzt nicht mehr der Fall. Es ist also notwendig die Selektionsprädikatmigration in die Joinreihenfolgebestimmung zu integrieren. Nur dann kann man auf gute Ergebnisse hoffen. Die Integration mit einem Ansatz des dynamischen Programmierens ist problematisch, da dort Lösungen verworfen werden, die unter Umständen zur Optimalen Lösung führen, wenn ein Selektionsprädikat nicht ganz nach unten durchgeschoben wird [356].

Eine Teillösung wird dort auch angedeutet. Ist der Rang eines Selektionsprädikates größer als jeder Rang jedes Plans einer Menge von Verbunden, so ist das Selektionsprädikat in einem optimalen Baum oberhalb all dieser Verbundoperationen plaziert. Ein entsprechender Algorithmus hat aber, wenn er beispielsweise nur Left-deep-trees erzeugt, eine Worst-case-Komplexität von $O(n^4 n!)$.

30.1.14 Sciore & Sieg

Die Hauptidee von Sciore und Sieg ist es, die Regelmenge in Module zu organisieren und jedem Modul eine eigene Suchstrategie, Kostenberechnung und Regelmenge zuzuordnen. Module können andere Module explizit aufrufen, oder implizit ihre Ausgabemenge an das nächste Modul weiterleiten.

30.1.15 Secondo

Gueting

30.1.16 Squirrel

Der erste Ansatz eines regelbasierten Optimierers, Squirrel, kann auf das Jahr 1975 zurückgeführt werden [714]. Man beachte, daß dieses Papier vier Jahre älter ist als das vielleicht am häufigsten zitierte Papier über den System R Optimierer [672], der jedoch nicht regelbasiert, sondern fest verdrahtet ist.

Abbildung 30.7 gibt einen Überblick über den Aufbau von Squirrel. Nach der syntaktischen Analyse liegt ein Operatorgraph vor. Dieser ist in Squirrel zunächst auf einen Operatorbaum beschränkt. Zur Behandlung von gemeinsamen Teilausdrücken wird das Anlegen von temporären Relationen, die den

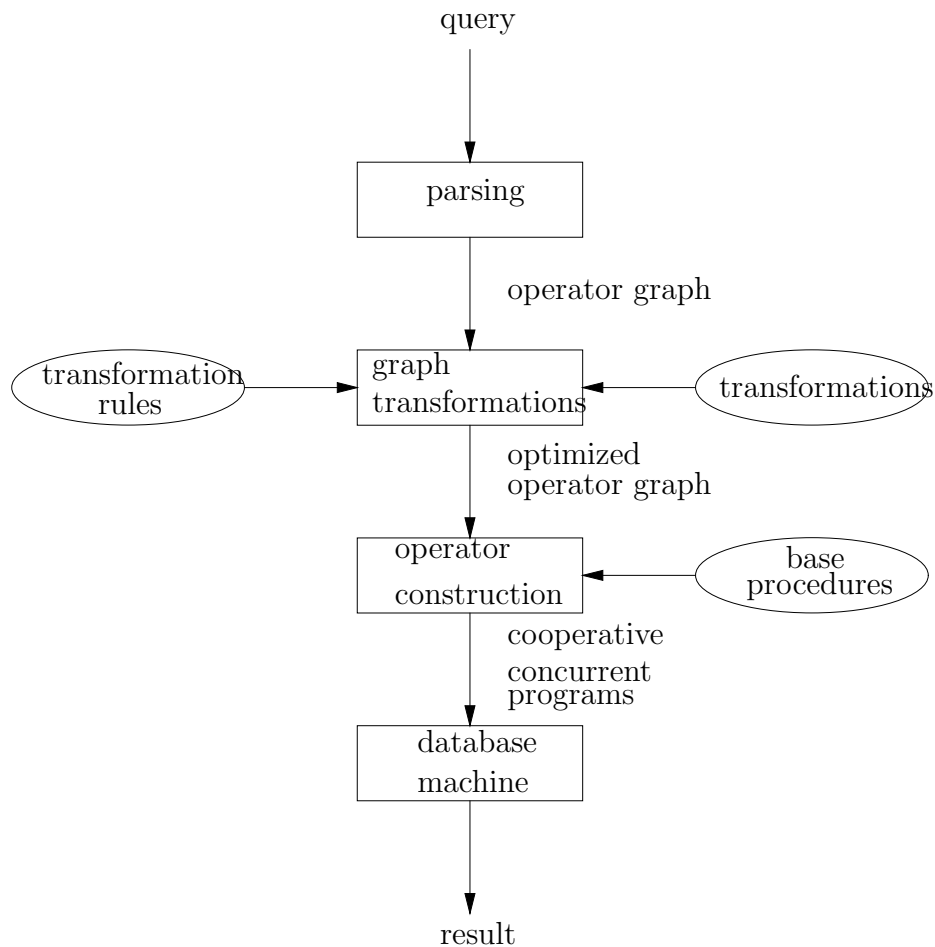


Figure 30.7: Die Squirrelarchitektur

gemeinsamen Teilausdrücken entsprechen, vorgeschlagen. Diese temporären Relationen ersetzen dann die gemeinsamen Teilausdrücke. Dadurch ist es möglich, sich auf Operatorbäume zu beschränken.

Der Operatorbaum wird dann in einen optimierten Operatorbaum transformiert. Hierzu werden Regeln, die den algebraischen Gleichungen entsprechen, verwendet. Die Anwendung dieser Transformationsregeln ist rein heuristisch gesteuert. Die Heuristik selber ist in den Transformationsanwendungsregeln abgelegt. Eine dieser Regeln sagt beispielsweise, daß Projektionen nur dann nach unten geschoben werden, wenn die Operation, über die die Projektion als nächstes geschoben werden soll, keine Verbundoperation ist. Neben den Standardregeln, die das Vertauschen von relationalen Operatoren ermöglichen, gibt es Regeln, die es erlauben, relationale Ausdrücke in komplexe boolesche Ausdrücke, die dann als Selektionsprädikate Verwendung finden, zu überführen. Dies ist der erste Vorschlag, nicht nur primitive Selektionsprädikate in Form von Literalen, sondern auch komplexere Ausdrücke mit booleschen Verknüpfungen zu verwenden. Auf die Optimierung dieser Ausdrücke wird jedoch nicht weiter

eingegangen.

Die wesentliche Aufgabe der Operatorkonstruktion ist die Auswahl der tatsächlichen Implementierungen der Operatoren im Operatorgraph unter optimaler Ausnutzung gegebener Sortierreihenfolgen. Auch diese Phase der Optimierung ist in Squirrel nicht kostenbasiert. Sie wird durch zwei Durchläufe durch den Operatorgraphen realisiert. Der erste Durchlauf berechnet von unten nach oben die möglichen Sortierungen, die ohne zusätzlichen Aufwand möglich sind, da beispielsweise Relationen schon sortiert sind, und vorhandene Sortierungen durch Operatoren nicht zerstört werden. Im zweiten Durchlauf, von oben nach unten, werden Umsortierungen nur dann vorgenommen, wenn keine der im ersten Durchlauf berechneten Sortierungen eine effiziente Implementierung des zu konvertierenden Operators erlaubt. Beide Durchläufe sind mit Regelsätzen spezifiziert. Es ist bemerkenswert, daß die Anzahl der Regeln, 32 für den Aufwärtspañ und 34 für den Abwärtspañ, die Anzahl der Regeln für die Transformationsphase (insgesamt 7 Regeln), bei weitem übertrifft. Auch die Komplexität der Regeln ist erheblich höher.

Beide für uns interessante Phasen, die Operatorgraphtransformation und Operatorkonstruktion, sind mit Regeln spezifiziert. Es ist jedoch in beiden Phasen kein Suchprozeß nötig, da die Regeln alle Fälle sehr gezielt auflisten und somit einen eindeutigen Entscheidungsbaum beschreiben. Eine noch minutiösere Unterscheidung für die Erzeugung von Ausführungsplänen in der Operatorkonstruktionsphase gibt es nur noch bei Yao [836]. Diese haben auch den Vorteil, durch Kostenrechnungen belegt zu sein.

Da die Regeln in ihren Prämissen die Heuristik ihrer Anwendung mit kodieren und keine eigene Suchfunktion zur Anwendung der Regeln existiert, ist die Erweiterbarkeit sehr schwierig. Das Fehlen jeglicher Kostenbewertung macht eine Evaluation der Alternativen unmöglich. Daher ist es auch schwer, die einzelnen Komponenten des Optimierers, nämlich die Regeln, zu bewerten, zumal der transformierende Ansatz gewählt wurde. Der Forderung nach Vorhersagbarkeit und stetiger Leistungsabfall wird in diesem Ansatz ebenfalls nicht nachgegangen.

30.1.17 System R and System R*

30.1.18 Starburst and DB2

Starburst [229, 332] liegt ein erweiterbares relationales Datenmodell zugrunde. Die Anfragebearbeitung ist wie in System R und System R* in die zwei Schritte Anfrageübersetzung und -ausführung zergliedert [333]. Wir interessieren uns für den ersten Schritt, die Anfrageübersetzung. Einen Überblick gibt Abbildung 30.8. Nach der standardmäßigen Zerteilung liegt die Anfrage in der internen Darstellung QGM (Query Graph Model) vor. QGM ist an die Anfragesprache Hydrogen (ähnlich SQL) von Starburst angelehnt. Der wichtigste Grundbaustein von QGM ist der *select*-Operator. Dieser enthält eine Projektionsliste und das Anfrageprädikat in Graphform. Die Knoten sind markiert und referenzieren (gespeicherte) Relationen oder weitere QGM-Operatoren. Die Markierung ist entweder ein Quantor (\forall, \exists) oder die Mengenerzeugermarkierung




Figure 30.8: Starburst Optimierer

(F). Knoten, die mit F markiert sind, tragen zur Erzeugung des Ergebnisses eines Operators bei, die Quantorenmarkierungen zu dessen Einschränkung. Die Kanten sind mit den Prädikaten markiert. Es ergeben sich also Schleifen für nur eine Relation betreffende Prädikate. Weitere Operatoren sind *insert*, *update*, *intersection*, *union* und *group-by*. Daneben wird die QGM-Repräsentation einer Anfrage mit Schemainformation und statistischen Daten angereichert. Sie dient also auch als Sammelbecken für alle die Anfrage betreffende Information.

Die QGM-Repräsentation dient der Anfragetransformation (Abb. 30.8) als Ausgangspunkt. Die Anfragetransformation generiert zu einer QGM-Repräsentation verschiedene äquivalente QGM-Repräsentationen. Die Anfragetransformation läßt sich, abgesehen von den Darstellungsunterschieden von QGM und Hydrogen, als eine Variante der Source-level-Transformationen ansehen. Sie wird regelbasiert implementiert, wobei C die Regelsprache ist. Eine Regel besteht aus 2 Teilen, einer Bedingung und einer Aktion. Jeder Teil wird durch eine C-Prozedur beschrieben. Dadurch erübrigt sich die Implementierung eines allgemeinen Regelinterpreters mit Pattern-matching. Regeln können in Gruppen zusammengefaßt werden. Der aktuelle Optimierer umfaßt drei Klassen von Regeln:

1. Migration von Prädikaten

2. Migration von Projektionen
3. Verschmelzung von Operationen

Für die Ausführung der Regeln stehen drei verschiedene Suchstrategien zur Verfügung:

1. sequentiell,
2. prioritätsgesteuert und
3. zufällig, gemäß einer gegebenen Verteilung.

Die Teilgraphen der QGM-Repräsentation, auf die Regeln anwendbar sind, können entweder durch eine depth-first oder eine breadth-first Suche bestimmt werden. Falls mehrere alternative QGM-Repräsentationen existieren (was meistens der Fall ist), wird ein *Choose*-Operator [307] verwendet, der die verschiedenen QGMs in einen QGM zusammenbaut. Die nachfolgende Phase wählt dann kostenbasiert einen dieser alternativen QGMs aus. Dies ist nicht zwingend, die Auswahl kann auch erst zur Auswertungszeit stattfinden. Begründet wird dieses Vorgehen damit, daß keine Kosten für QGMs berechnet werden können, und somit keine Bewertung eines QGMs stattfinden kann. Wie die Autoren selbst anmerken, ist dieser Umstand sehr mißlich, da keine Alternativen verworfen werden können. Sie kündigen daher Untersuchungen an, die Transformation (Schritt 2) mit der Planoptimierung (Schritt 3) zu verschmelzen. Um eine gewisse Kontrolle über das Verhalten der Transformation zu haben, kann diesem Schritt ein "budget" mitgegeben werden, nach dessen Ablauf der Schritt beendet wird. Die genaue Funktionsweise des "budget" ist leider nicht erläutert.

Der Schritt der Planoptimierung (s. Abb. 30.8) kann mit der bisherigen Optimierung verglichen werden. Sie arbeitet regelbasiert, benutzt aber nicht den transformierenden, sondern den generierenden Ansatz [488]. Aus Basisoperationen – LOLEPOPs (LOW-LEvel Plan OPERator) genannt – werden mit (grammatischen) Regeln – STARS (strategy alternative rules) genannt – (alternative) Auswertungspläne erzeugt. LOLEPOPs entstammen der um SCAN, SORT und ähnliche physische Operatoren angereicherten relationalen Algebra. Ein Auswertungsplan ist dann ein Ausdruck von geschachtelten Funktionsaufrufen, wobei die Funktionen den LOLEPOPs entsprechen.

Ein STAR definiert ein benanntes parametrisiertes Objekt, das einem Nicht-terminalsymbol entspricht. Er besteht aus einer Menge von alternativen Definitionen, die jede aus einer Bedingung für die Anwendbarkeit und der Definition eines Plans bestehen. Der generierte Plan kann LOLEPOPs (entsprechen Terminalsymbolen) und STARS referenzieren. Ein rootSTAR entspricht dem Startsymbol der Grammatik. STARS ähneln den Regeln, die in Genesis nicht nur für den Optimierer, sondern für das ganze DBMS eingesetzt werden, um alternative Implementierungen zu erhalten [52, 50, 54, 53, 55]. Um erzeugte Alternativen für einen Plan zusammensetzen und zu verhindern, daß diese Alternativen die Anzahl der Pläne in denen diese vorkommen, vervielfachen, wird ein Glue-Mechanismus eingesetzt. Dieser hat den *Choose*-Operator als Wurzel. Darunter hängen dann Alternativen, die beispielsweise einen Strom

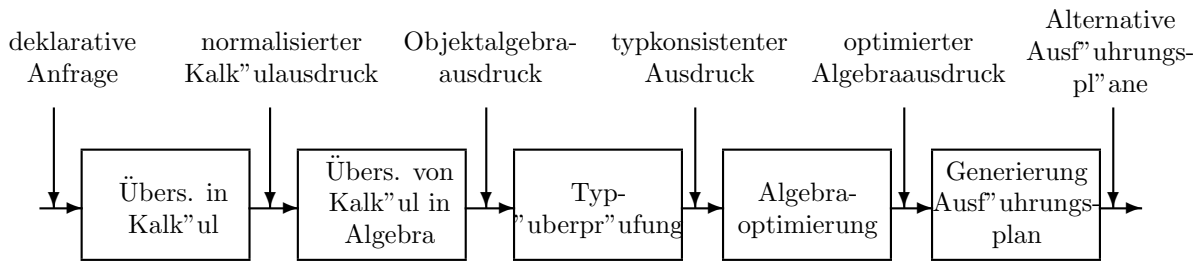


Figure 30.9: Der Optimierer von Straube

mit gewissen Eigenschaften (Sortierung, Lokation) erzeugen. Von diesen Alternativen werden nur diejenigen betrachtet, die die geringsten Kosten bei gleichen Eigenschaften haben [463]. Die Kosten beziehen sich dabei immer nur auf den bisher erreichten Teilplan.

Der Aufbau eines Auswertungsplanes erfolgt Bottom-up. Die Menge der anwendbaren STARS wird in einer ToDo-Liste gehalten. Diese ist eine sortierte Liste. Hiermit können dann verschiedene Suchstrategien implementiert werden, indem verschiedene Sortierungen für die ToDo-Liste Verwendung finden [463]. Ein Vorteil des STAR-Ansatzes ist die Vermeidung von Pattern-matching. Dies erlaubt es, die STARS zu interpretieren [463].

Die Beurteilung der Erweiterbarkeit ist sehr schwierig. Zum einen handelt es sich um einen erweiterbaren Optimierer, da sowohl LOLEPOPs als auch STARS hinzugefügt werden können. Der Glue-Mechanismus kann ebenfalls spezifiziert werden, ohne in die Implementierung einzugreifen. Das Problem ist lediglich die Komplexität dieser Änderungen. Man kann diesen Ansatz daher vielleicht als bedingt erweiterbar kennzeichnen.

Eine Trennung der Optimierung in verschiedene Phasen wirft die damit verbundenen Probleme auf. Wie oben bereits angeführt, kündigen die Autoren weitere Untersuchungen an, um eine Verschmelzung der Phasen zu ermöglichen. Da keine Alternativen verworfen werden, ist es potentiell möglich, den optimalen Auswertungsplan zu errechnen. Schwer zu sehen ist jedoch, wie ein stetiger Leistungsabfall zu realisieren ist. Gleiches gilt für die Evaluierbarkeit der einzelnen Komponenten (STARS).

More on Starburst can be found in [586, 587].

30.1.19 Der Optimierer von Straube

In seiner Dissertation stellt Straube den von ihm entwickelten Optimierer dar [739]. Die Ergebnisse dieser Arbeit flossen in eine Reihe von Veröffentlichungen ein [769, 736, 737, 735, 738]. Der Aufbau des Optimierers ist in Abbildung 30.9 skizziert. Eine Anfrage wird zunächst in den Objektkalkül übersetzt und von dort in die Objektalgebra. Hier findet dann zunächst eine Typüberprüfung statt. Danach beginnt die eigentliche Optimierung. Diese besteht aus zwei Phasen, der algebraischen Optimierung und der Generierung des Ausführungsplans.

Die erste Phase, die algebraische Optimierung, folgt dem Transformationssparadigma. Die algebraischen Ausdrücke werden mit Hilfe von Regeln in

äquivalente algebraische Ausdrücke transformiert. Straube beschränkt sich dabei im wesentlichen auf die Formulierung der Regeln. Für die Abarbeitung der Regeln schlägt er lediglich die Verwendung des Exodus-Optimierergenerators [303] oder des Anfrageumformers von Starburst [352] vor.

Die zweite Phase, die Generierung der Ausführungspläne, ist nicht regelbasiert. Ihr liegt eine sogenannte Ausführungsplanschablone zu Grunde. Sie ist vergleichbar mit einem Und/Oder-Baum, der alle möglichen Ausführungspläne implizit beinhaltet. Zur Generierung eines konkreten Ausführungsplans wird der durch die Ausführungsschablone aufgespannte Suchraum vollständig durchsucht. Der billigste Ausführungsplan kommt dann zur Abarbeitung.

Da ein regelbasierten Ansatz für die erste Phase gewählt wurde und die Verwendung des Exodus-Optimierergenerators oder des Starburst-Anfrageumformers vorgeschlagen wird, verweisen wir für die Bewertung dieser Phase auf die entsprechenden Abschnitte.

Die zweite Phase ist voll auskodiert und damit schlecht erweiterbar. Eine frühe Bewertung der Alternativen ist nicht ausgeschlossen, wird aber nicht vorgenommen. Ein vollständiges Durchsuchen verhindert natürlich auch den stetigen Leistungsabfall des Optimierers.

Erschwerend für den gewählten Ansatz kommt die Zweiphasigkeit hinzu. Es ist schwierig zu sehen, wie eine phasenübergreifende Kontrolle auszusehen hat, die zumindest potentiell Optimalität gewährleistet. Die Evaluierbarkeit einzelner Komponenten des Optimierers ist nicht möglich.

30.1.20 Other Query Optimizer

Neben den in den vorangehenden Abschnitten erwähnten Optimierern gibt es noch eine ganze Reihe anderer, die aber nicht im einzelnen vorgestellt werden sollen. Erwähnt werden sollen noch die Systeme Probe [191, 190, 555] und Prima [348, 346]. Der Schwerpunkt bei der Optimierung liegt im Primasystem auf dem dynamischen Zusammenbau von Molekülen. Es wäre zu untersuchen, ob ein Assembly-Operator (s. [412]) hier von Nutzen wäre. Besonders erwähnenswert ist noch eine Arbeit, die Optimierungsmöglichkeiten für die Datenbankprogrammiersprache FAD vorstellt [774]. Diese Arbeit stellt einen ersten Schritt in Richtung eines Optimierers für eine General-purpose-Programmiersprache dar. Ein wesentlicher Punkt ist dabei, daß auch ändernde Operationen optimiert werden. Der Optimierer ist in zwei Module (RWR and OPT) eingeteilt. RWR ist ein Sprachmodul, das die Übersetzung von FAD in ein internes FAD vornimmt. Immer wenn RWR einen Ausdruck erkennt, der in der vom Optimierer OPT bearbeitbaren Sprache ausgedrückt werden kann, so wird dieser an den Optimierer weitergegeben und dort optimiert. Es wird Exhaustive search als Suchstrategie für den Optimierer vorgeschlagen.

Im erweiterten O_2 -Kontext wurde das Zerteilen von Pfadausdrücken weiter untersucht [159]. Es werden die Vorteile einer typisierten Algebra für diese Zwecke herausgearbeitet. Eine graphische Notation sorgt für eine anschauliche Darstellung. Ihr besonderes Augenmerk richten die Autoren auf die Faktorisierung gemeinsamer Teilausdrücke. Einige der Ersetzungsregeln sind aus [693] und [692] entnommen und werden gewinnbringend eingesetzt, so beispielsweise die

Ersetzung von Selektionen durch Verbundoperatoren.

Ebenfalls erwähnt wurden bereits die Arbeiten im Orion-Kontext [45, 46, 423], die sich auf die Behandlung von Pfadausdrücken konzentrieren. Auch hier wurde ein funktionsfähiger Optimierer entwickelt.

Wie bereits erwähnt, stammt der erste regelbasierte Optimierer von Smith und Chang [714]. Doch erst die neueren Arbeiten führten zu einer Blüte des regelbasierten Ansatzes. Hier ist insbesondere die Arbeit von Freytag zu erwähnen, die diese Blüte mit initiierte [251]. Dort wird gezeigt, wie man mit Hilfe eines regelbasierten Ansatzes den Optimierer von System R [672] emulieren kann. Die Eingabe besteht aus einem Lisp-ähnlichen Ausdruck:

```
(select <proj-list>
      <sel-pred-list>
      <join-pred-list>
      <table-list>)
```

Die Projektionsliste besteht aus Attributspezifikationen der Form

```
<rel-name>.<attr-name>
```

Diese werden auch für die Selektionsprädikate und Joinprädikate verwendet. Die Algebra beinhaltet sowohl Operatoren der logischen als auch der physischen Algebra. Im einzelnen gibt es Scan-, Sort-, Projektions-, Verbundoperatoren in einer logischen und verschiedenen physischen Ausprägungen. Die Erzeugung der Auswertungspläne wird in verschiedene Schritte unterteilt, die wiederum in Teilschritte zerlegt sind (siehe Abb. 30.3).

Zunächst erfolgt die Übersetzung in die logische Algebra. Hier werden Scan-Operatoren um die Relationen gebaut und Selektionen, die nur eine Relation betreffen, in die Scan-Operatoren eingebaut. Der zweite Schritt generiert Zugriffspläne, indem der Scan-Operator durch einen einfachen File-Scan (FSCAN) ersetzt wird oder falls möglich, durch einen Index-Scan (ISCAN). Der dritte Schritt generiert zunächst verschiedene Verbundreihenfolgen und bestimmt anschließend die Verbundmethoden. Wie in System R wird zwischen dem Sortiere- und Mische-Verbund und dem Verbund durch geschachtelte Schleifen unterschieden. Es werden keinerlei Aussagen über die Auswahl einer Suchstrategie gemacht. Ziel ist es vielmehr, durch die Modellierung des System R Optimierers mit Hilfe eines Regelsystems die prinzipielle Brauchbarkeit des regelbasierten Ansatzes nachzuweisen.

Man beachte auch die erwähnte Arbeit von Sciore und Sieg zur Modularisierung von regelbasierten Optimierern [668]. Die Hauptidee von Sciore und Sieg ist es, die Regelmenge in Module zu organisieren und jedem Modul eine eigene Suchstrategie, Kostenberechnung und Regelmenge zuzuordnen. Module können andere Module explizit aufrufen oder implizit ihre Ausgabemenge an das nächste Modul weiterleiten. Der erste Optimierer des GOM-Systems ist ebenfalls regelbasiert [415, 414]. Die gesamte Regelmenge wurde hier in Teilmengen ähnlich zu den Modulen organisiert. Die Steuerung zwischen den Teilmengen erfolgt durch ein heuristisches Netz, das angibt in welchen Fällen zu welcher weiteren Teilmenge von Regeln zu verzweigen ist. Die Strukturierung des Optimiererwissens steht auch in [498] im Vordergrund.

In diesem Zusammenhang, der Strukturierung von Optimierern und der Wiederverwendbarkeit einzelner Teile, sei noch einmal ausdrücklich auf die Arbeiten von Batory [53] aus dem Genesiskontext hingewiesen (s. auch Abschnitt 30.1.7). Der dort leider ein wenig zu kurz kommende Aspekt der Wiederverwendbarkeit von Suchfunktionen wird in einer Arbeit von Lanzelotte und Valduries [452] ausführlicher behandelt. Hier wurde eine Typhierarchie existierender Suchfunktionen entworfen und deren Schnittstellen vereinheitlicht. Die Suchfunktionen selbst wurden modularisiert. Weitere Arbeiten aus derselben Gruppe beschäftigen sich mit der Optimierung von objektorientierten Anfragen [451, 455], wobei hier die Behandlung von Pfaden im Vordergrund steht. Eine neuere Arbeit beschäftigt sich mit der Optimierung von rekursiven Anfragen im objektorientierten Kontext [453].

Viele kommerzielle Systeme besitzen eine Anfragesprache und einen Optimierer. Einer der wenigen Optimierer, die auch in der Literatur beschrieben werden, ist der von ObjectStore [554]. Durch die einfache Anfragesprache, die nur Teilmengenbestimmung erlaubt, und die strikte Verwendung von *C*-Semantik für boolesche Ausdrücke sind die meisten Optimierungsmöglichkeiten jedoch ausgeschlossen, und der “Optimierer” ist daher sehr einfach.

30.2 Commercial Query Compiler

30.2.1 The DB 2 Query Compiler

30.2.2 The Oracle Query Compiler

Oracle still provides two modes for its optimizer. Dependent on the user specified optimizer mode, a query is optimized either by the rule-based optimizer (RBO) or by the cost-based optimizer (CBO). The RBO is a heuristic optimizer that resembles the simple optimizer of chapter 2. Here we concentrate on the more powerful CBO. The user can also determine whether the optimizer should optimize for throughput or response time.

- nested loop join, nested loop outer join, index nested loop joins, sort merge join, sort merge outer join, hash joins, hash outer join, cartesian join, full outer join, cluster join, anti-joins, semi-joins, uses bitmap indexes for star queries
- sort group-by,
- bitmap indexes, bitmap join indexes
- index skip scans
- partitioned tables and indexes
- index-organized tables
- reverse key indexes
- function-based indexes

- SAMPLE clause in SELECT statement
- parallel query and parallel DML
- star transformations and star joins
- query rewrite with materialized views
- cost: considers CPU, I/O, memory
- access path: table scan, fast full index scan, index scan, ROWID scans (access ROW by ROWID), cluster scans, hash scans. [former two with prefetching] index scans:
 - index unique scan (UNIQUE or PRIMARY KEY constraints)
 - index range scan (one or more leading columns or key)
 - index range scan descending
 - index skip scan (> 1 leading key values not given)
 - index full scan, index fast full scan
 - index joins (joins indexes with hash join, resembles index anding)
 - bitmap joins (index anding/oring)
 - cluster scan: for indexed cluster to retrieve rows with the same cluster id
 - hash scan: to locate rows in a hash cluster

CBO: parsed quer \rightarrow [query transformer] \rightarrow [estimator] \rightarrow [plan generator] 1-16.

after parser: nested query blocks

simple rewrites:

- eliminate between
- eliminate x in (c1 ... cn) (also uses IN-LIST iterator as outer table constructor in a d-join or nested-loop join like operation.

query transformer:

- view merging
- predicate pushing
- subquery unnesting
- query rewrite using materialized views (cost based)

remaining subplans for nested query blocks are ordered in an efficient manner
plan generator:

- choose access path, join order (upper limit on number of permutations considered), join method.

- generate subplan for every block in a bottom-up fashion
- (> 1 for still nested queries and unmerged views)
- stop generating more plans if there already exists a cheap plan
- starting plan: order by their effective cardinality
- considers normally only left-deep (zig-zag) trees.
- single row joins are placed first (based on unique and key constraints.
- join statement with outer join: table with outer join operator must come after the other table in the condition in the join order. optimizer does not consider join orders that violate this rule.
- NOT IN (SELECT ...) becomes a anti-join that is executed as a nested-loop join by default unless hints are given and various conditions are met which allow the transformation of the NOT IN uncorrelated subquery into a sort-merge or hash anti-join.
- EXISTS (SELECT ...) becomes a semi-join. execution as index nested loops, if there is an index. otherwise a nested-loop join is used by default for EXISTS and IN subqueries that cannot be merged with the containing query unless a hint specifies otherwise and conditions are met to allow the transformation of the subquery into a sort-merge or hash semi-join.

- star query detection

cost:

- takes unique/key constraints into consideration
- low/high values and uniform distribution
- host variables: guess small selectivity value to favor index access
- histograms
- common subexpression optimization
- complex view merging
- push-join predicate
- bitmap access paths for tables with only B-tree indexes
- subquery unnesting
- index joins

rest:

- Oracle allows user hints in SQL statements to influence the Optimizer. for example join methods can be given explicitly

parameters:

- HASH_AREA_SIZE
- SORT_AREA_SIZE
- DB_FILE_MULTIBLOCK_READ_COUNT (number of prefetched pages)

statistics:

- table statistics
number of rows, number of blocks, average row length
- column statistics
number of distinct values, number of nulls, data distribution
- index statistics
number of keys, (from column statistics?) number of leaf blocks, levels, clustering factor (collocation amount of the index block/data blocks, 3-17)
- system statistics
I/O performance and utilization, cpu performance and utilization

generating statistics:

- estimation based on random data sampling
(row sampling, block sampling)
- exact computation
- user-defined statistics collection methods

histograms:

- height-based histograms (approx. equal number of values per bucket)
- value-based histograms
used for number of distinct values \leq number of buckets
- support of index-only queries
- index-organized tables
- bitmap indexes (auch fuer null-werte `x <> const`)
- convert b-tree result RID lists to bitmaps for further bitmap anding
- bitmaps and count
- bitmap join index
- cluster tables (cluster rows of different tables on the same block)
- hash clusters

- hint: USE_CONCAT: OR ::= UNION ALL
- hint: STAR_TRANSFORMATION: see Oracle9i Database Concepts
- NOT IN ::= anti-join
- EXISTS ::= special join preserving duplicates and adding no phantom duplicates (semi-join) (5-27)
- continue 5-35

30.2.3 The SQL Server Query Compiler

Part VI
Selected Topics

Chapter 31

Generating Plans for Top-N-Queries?

31.1 Motivation and Introduction

motivation:

- first by user (ordered)
- optimize for n rows (user/cursor)
- exist(subquery) optimize for 1 row
- having count(*) <= n

31.2 Optimizing for the First Tuple

31.3 Optimizing for the First N Tuples

- nl-join instead of sm/hash join
- index access over table scan
- disable prefetching

[104, 105, 106] [124, 211] [230, 231, 372] [477]

[324] (also contains inverted list algorithms under frequent updates)

[477]

Chapter 32

Recursive Queries

Chapter 33

Issues Introduced by OQL

33.1 Type-Based Rewriting and Pointer Chasing Elimination

The first rewrite technique especially tailored for the object-oriented context is *type-based rewriting*. Consider the query

```
select distinct sn, ssn, ssa
from          s in Student
```

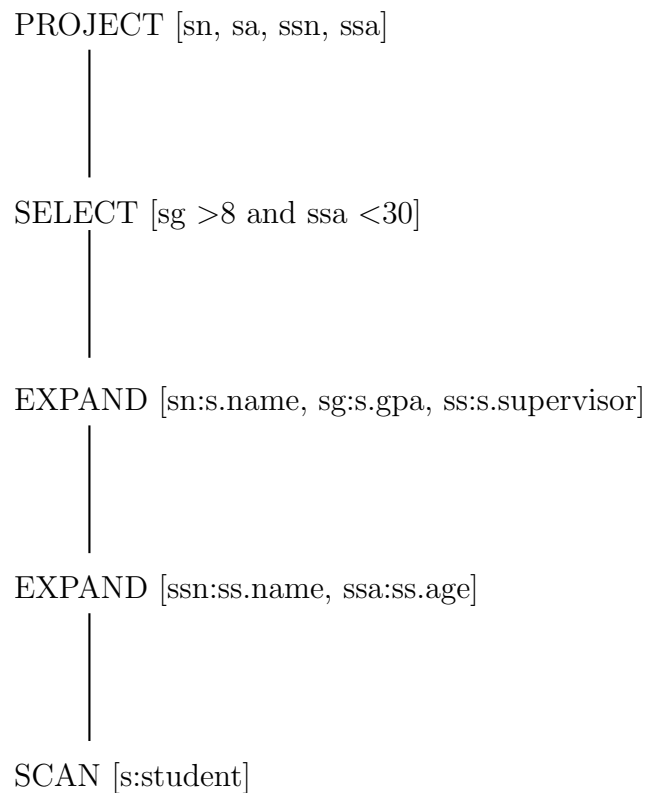


Figure 33.1: Algebraic representation of a query


```

where          sg > 8 and ssa < 30
define         sn = s.name
                 sg = s.gpa
                 ss = s.supervisor
                 ssn= ss.name
                 ssa= ss.age

```

The algebraic expression in Fig. 33.1 implies a scan of all students and a subsequent dereferencing of the *supervisor* attribute in order to access the supervisors. If not all supervisors fit into main memory, this may result in many page accesses. Further, if there exists an index on the supervisor's *age*, and the selection condition $ssa < 30$ is highly selective, the index should be applied in order to retrieve only those supervisors required for answering the query. Type-based rewriting enables this kind of optimization. For any expression of certain type with an associated extent, the extent is introduced in the *from* clause. For our query this results in

```

select distinct sn, pn, pa
from             s in Student, p in Professor
where           sg > 8 and pa < 30 and ss = p
define         sn = s.name
                 sg = s.gpa
                 ss = s.supervisor
                 pn = ss.name
                 pa = ss.age

```

As a side-effect, the attribute traversal from students via supervisor to professor is replaced by a join. Now, join-ordering allows for several new plans that could not be investigated otherwise. For example, we could exploit the above mentioned index to retrieve the young professors and join them with the students having a *gpa* greater than 8. The according plan is given in Fig. 33.2. Turning implicit joins or pointer chasing into explicit joins which can be freely reordered is an original query optimization technique for object-oriented queries. Note that the plan generation component is still allowed to turn the explicit join into an implicit join again.

Consider the query

```

select distinct p
from           p in Professor
where         p.room.number = 209

```

Straight forward evaluation of this query would scan all professors. For every professor, the *room* relationship would be traversed to find the room where the professor resides. Last, the room's number would be retrieved and tested to be 209. Using the *inverse relationship*, the query could as well be rewritten to

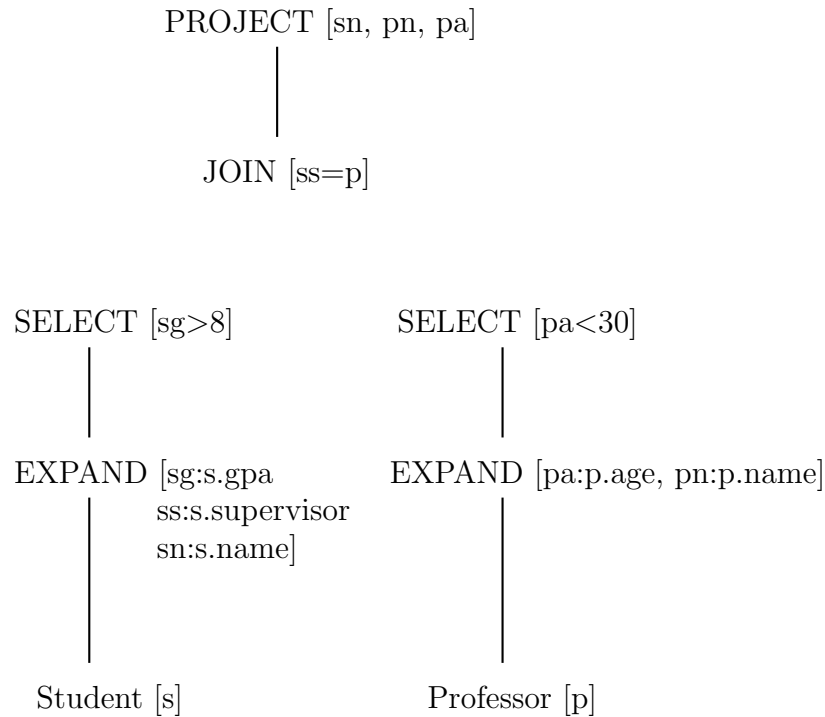


Figure 33.2: A join replacing pointer chasing

```

select distinct r.occupiedBy
from           r in Room
where         r.number = 209
  
```

The evaluation of this query can be much more efficient, especially if there exists an index on the room number. Rewriting queries by exploiting inverse relationships is another rewrite technique to be applied during Rewrite Phase I.

33.2 Class Hierarchies

Another set of equivalences known from the relational context involves the UNION operator (\cup) and plays a vital role in dealing with class/extent hierarchies. Consider the simple class hierarchy given in Figure 33.3. Obviously, for the user, it must appear that the extent of *Employee* contains all *Managers*. However, the system has different alternatives to implement extents. Most OBMSs organize an object base into areas or volumes. Each area or volume is then further organized into several files. A file is a logical grouping of objects not necessarily consisting of subsequent physical pages on disk. Files don't share pages.

The simplest possible implementation to scan all objects belonging to a certain extent is to perform an area scan and select those objects belonging to the extent in question. Obviously, this is far too expensive. Therefore, some more so-

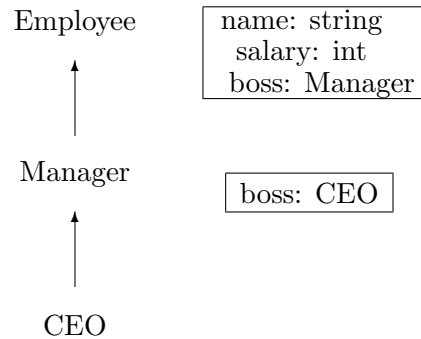


Figure 33.3: A Sample Class Hierarchy

phisticated possibilities to realize extents and scans over them are needed. The different possible implementations can be classified along two dimensions. The first dimension distinguishes between logical and physical extents, the second distinguishes between strict and (non-strict) extents.

Logical *vs.* Physical Extents

An extent can be realized as a collection of object identifiers. A scan over the extent is then implemented by a scan over all the object identifiers contained in the collection. Subsequently, the object identifiers are dereferenced to yield the objects themselves. This approach leads to logical extents. Another possibility is to implement extent membership by physical containment. The best alternative is to store all objects of an extent in a file. This results in physical extents. A scan over a physical extent is then implemented by a file scan.

Extents *vs.* Strict Extents

A strict extent contains the objects (or their OIDs) of a class excluding those of its subclasses. A non-strict extent contains the objects of a class and all objects of its subclasses.

Given a class C , any strict extent of a subclass C' of C is called a subextent of C .

Obviously, the two classifications are orthogonal. Applying them both results in the four possibilities presented graphically in Fig. 33.4. [166] strongly argues that strict extents are the method of choice. The reason is that only this way the query optimizer might exploit differences for extents. For example, there might be an index on the *age* of *Manager* but not for *Employee*. This difference can only be exploited for a query including a restriction on *age*, if we have strict extents.

However, strict extents result in initial query plans including UNION operators. Consider the query

```

select  e
from    e in Employee
where   e.salary > 100.000
  
```


Chapter 34

Issues Introduced by XPath

34.1 A Naive XPath-Interpreter and its Problems

34.2 Dynamic Programming and Memoization

[285, 287, 286]

34.3 Naive Translation of XPath to Algebra

34.4 Pushing Duplicate Elimination

34.5 Avoiding Duplicate Work

34.6 Avoiding Duplicate Generation

[357]

34.7 Index Usage and Materialized Views

[43]

34.8 Cardinalities and Costs

34.9 Bibliography

Chapter 35

Issues Introduced by XQuery

35.1 Reordering in Ordered Context

35.2 Result Construction

[242, 243] [687]

35.3 Unnesting Nested XQueries

Unnesting with error: [572]
[508, 510, 509, 511]

35.4 Cardinalities and Cost Functions

cardinality: [142, 818, 819, 651] [7]
XPathLearner: [483]
Polyzotis et al (XSKETCH): [592, 593, 590], [594]

35.5 Bibliography

[512] [765] [199]
Numbering: [241] Timber [392] TAX Algebra [393], physical algebra of Timber [573]
Structural Joins [21, 716]
SAL: [63], TAX: [393], XAL: [249]

- XML Statistics for hidden web: [8]
- XPath selectivity for internet scale: [7]
- StatiX: [250]
- IMAX: incremental statistics [606]
- Metrics for XML Document Collections: [428]

- output size containment join: [788]
- Bloom Histogram: [789]

View and XML: [2]

Quilt: [116]

Timber: [392] Monet: [663] Natix: NoK: [845]

Correlated XPath: [846]

Wood: [810, 811, 812]

Path based approach to Storage (XRel): [839]

Grust: [317, 319, 318, 320, 758]

Liefke: Loop fusion etc.: [482]

Benchmarking: XMach-1: [83], MBench: [641] XBenck: [566, 831, 832],

XMark: [664] XOO7: [91]

Rewriting: [203, 309, 310]

[202, 368]

Incremental Schema Validation: [87, 571]

Franklin (filtering): [206]

Chapter 36

Outlook

What we did not talk about: multiple query optimization, semantic query optimization, special techniques for optimization in OBMSs, multi-media data bases, object-relational databases, spatial databases, temporal databases, and query optimization for parallel and distributed database systems.

Multi Query Optimization? [675]

Parametric/Dynamic/Adaptive Query Optimization? [30, 31, 32, 28, 307, 300] [388, 389, 405, 770]
[40]

Parallel Database Systems?

Distributed Database Systems? [439]

Recursive Queries?

Multi Database Systems?

Temporal Database Systems?

Spatial Database Systems?

Translation of Triggers and Updates?

Online Queries (Streams)?

Approximate Answers? [278]

Appendix A

Query Languages?

A.1 Designing a query language

requirements

design principles for object-oriented query languages: [367] [74]

A.2 SQL

A.3 OQL

A.4 XPath

A.5 XQuery

A.6 Datalog

Appendix B

Query Execution Engine (?)

- Overview Books: [347, 270]
- Overview: Graefe [294, 295]
- Implementation of Division [291, 299, 301]
- Implementation of Division and set-containment joins [603]
- Hash vs. Sort: [296, 304]
- Heap-Filter Merge Join: [293]
- Hash-Teams

Appendix C

Glossary of Rewrite and Optimization Techniques

trivopt Triviale Auswertungen bspw. solche für widersprüchliche Prädikate werden sofort vorgenommen. Dies ist eine Optimierungstechnik, die oft bereits auf der Quellebene durchgeführt wird.

pareval Falls ein Glied einer Konjunktion zu *false* evaluiert, werden die restlichen Glieder nicht mehr evaluiert. Dies ergibt sich automatisch durch die Verwendung von hintereinanderausgeführten Selektionen.

pushnot Falls ein Prädikat die Form $\neg(p_1 \wedge p_2)$ hat, so ist *pareval* nicht anwendbar. Daher werden Negationen nach innen gezogen. Auf $\neg p_1 \vee \neg p_2$ ist *pareval* dann wieder anwendbar. Das Durchschieben von Negationen ist auch im Kontext von NULL-Werten unabdingbar für die Korrektheit. Dies ist eine Optimierungstechnik, die oft bereits auf der Quellebene durchgeführt wird.

bxp Verallgemeinert man die in *pareval* und *notpush* angesprochene Problematik, so führt dies auf die Optimierung von allgemeinen booleschen Prädikaten.

trans Durch Ausnutzen der Transitivität von Vergleichsoperationen können neue Selektionsprädikate gewonnen und Konstanten propagiert werden. Diese Optimierungstechnik erweitert den Suchraum und wird ebenfalls auf der Quellebene durchgeführt. Bei manchen Systemen wird dieser Schritt nicht durchgeführt, falls sehr viele Relationen zu joinen sind, um den Suchraum nicht noch weiter zu vergrößern [275, 276].

selpush Selektionen werden so früh wie möglich durchgeführt. Diese Technik führt nicht immer zu optimalen Auswertungsplänen und stellt somit eine Heuristik dar. Diese Optimierungstechnik schränkt den Suchraum ein.

projpush Die Technik zur Behandlung von Projektionen ist nicht ganz so einfach wie die der Selektion. Zu unterscheiden ist hier, ob es sich um eine Projektion mit Duplikatelimination handelt oder nicht. Je nach dem

ist es sinnvoll, die Projektion zur Wurzel des Operatorgraphen zu verschieben oder zu den Blättern hin. Die Projektion verringert den Speicherbedarf von Zwischenergebnissen, da die Tupel weniger Attribute enthalten. Handelt es sich um eine duplikateliminiierende Projektion, so wird möglicherweise auch die Anzahl der Tupel verringert. Duplikatelimination als solche ist aber eine sehr teure Operation. Diese wird üblicherweise durch Sortieren implementiert. Bei großen Datenmengen gibt es allerdings bessere Alternativen. Auch Hash-basierte Verfahren eignen sich zur Duplikateliminiierung. Diese Optimierungstechnik schränkt den Suchraum ein.

grouppush Pushing a grouping operation past a join can lead to better plans.

crossjoin Ein Kreuzprodukt, das von einer Selektion gefolgt wird, wird wenn immer möglich in eine Verbundoperation umgewandelt. Diese Optimierungstechnik schränkt den Suchraum ein, da Pläne mit Kreuzprodukten vermieden werden.

nocross Kreuzprodukte werden wenn immer möglich vermieden oder, wenn dies nicht möglich ist, erst so spät wie möglich durchgeführt. Diese Technik verringert den Suchraum, führt aber nicht immer zu optimalen Auswertungsplänen.

semjoin Eine Verbundoperation kann durch eine Semiverbundoperation ersetzt werden, wenn nur die Attribute einer Relation weitere Verwendung finden.

joinor Die Auswertungsreihenfolge von Verbundoperationen ist kritisch. Daher wurden eine Reihe von Verfahren entwickelt, die optimale oder quasi-optimale Reihenfolge von Verbundoperationen zu bestimmen. Oft wird dabei der Suchraum auf Listen von Verbundoperationen beschränkt. Die Motivation hierbei ist das Verkleinern des Suchraums und die Beschränkung auf nur eine zu erzeugenden Zwischenrelation. Dieses Verfahren garantiert nicht mehr ein optimales Ergebnis.

joinpush Tables that are guaranteed to produce a single tuple are always pushed to be joined first. This reduces the search space. The single tuple condition can be evaluated by determining whether all key attributes of a relation are fully qualified. [275, 276].

elimredjoin Eliminate redundant join operations. See Sections... XXX

indnest Eine direkte Evaluierung von geschachtelten Anfragen wird durch geschachtelte Schleifen vorgenommen. Dabei wird eine Unteranfrage für jede erzielte Bindung der äußeren Anfrage evaluiert. Dies erfordert quadratischen Aufwand und ist deshalb sehr ineffizient. Falls die innere Anfrage unabhängig von der äußeren Anfrage evaluiert werden kann, so wird diese herausgezogen und getrennt evaluiert. Weitere Optimierungen geschachtelter Anfragen sind möglich.

unnest Entschachtelung von Anfragen [161, 162, 268, 424, 429, 430, 586, 719, 721, 722]

compop Oft ist es sinnvoll, mehrere Operationen zu einer komplexeren zusammenzufassen. Beispielsweise können zwei hintereinander ausgeführte Selektionen durch eine Selektion mit einem komplexeren Prädikat ersetzt werden. Ebenso kann auch das Zusammenfassen von Verbundoperationen, Selektionen und Projektionen sinnvoll sein.

comsubexpr Gemeinsame Teilausdrücke werden nur einfach evaluiert. Hierunter fallen zum einen Techniken, die das mehrmalige Lesen vom Hintergrundspeicher verhindern, und zum anderen Techniken, die Zwischenergebnisse von Teilausdrücken materialisieren. Letzteres sollte nur dann angewendet werden, falls die k-malige Auswertung teurer ist als das einmalige Auswerten und das Erzeugen des Ergebnisses mit k-maligem Lesen, wobei k die Anzahl der Vorkommen im Plan ist.

dynminmax Dynamisch gewonnene Minima und Maxima von Attributwerten können für die Erzeugung von zusätzlichen Restriktionen herangezogen werden. Diese Technik funktioniert auch sehr gut für unkorrelierte Anfragen. Dabei werden min- und max-Werte herangezogen um zusätzliche Restriktionen für die Anfrage zu gewinnen. [429, 275, 276]

pma Predicate Move around moves predicates between queries and subqueries. Mostly they are duplicated in order to yield as many restrictions in a block as possible [474]. As a special case, predicates will be pushed into view definitions if they have to be materialized temporarily [275, 276].

exproj For subqueries with exist prune unnecessary entries in the **select** clause. The intention behind is that attributes projected unnecessarily might influence the optimizer's decision on the optimal access path [275, 276].

vm View merging expands the view definition within the query such that it can be optimized together with the query. Thereby, duplicate accesses to the view are resolved by different copies of the views definition in order to facilitate unnesting [275, 276, 586].

inConstSet2Or A predicate of the form $x \in \{a_1, \dots, a_n\}$ is transformed into a sequence of disjunctions $x = a_1 \vee \dots \vee x = a_n$ if the a_i are constants in order to allow index or-ing (TID list operations or bitvector operations) [275, 276].

like1 If the like predicate does not start with %, then a prefix index can be used.

like2 The pattern is analyzed to see whether a range of values can be extracted such that the pattern does not have to be evaluated on all tuples. The result is either a pretest or an index access. [275, 276].

like3 Special indexes supporting like predicates are introduced.

sort Vorhandene Sortierungen können für verschiedene Operatoren ausgenutzt werden. Falls keine Sortierung vorhanden ist, kann es sinnvoll sein, diese zu erzeugen [709]. Z.B. aufeinanderfolgende joins, joins und gruppierungen. Dabei kann man die Gruppierungsattribute permutieren, um sie mit einer gegebenen Sortierreihenfolge in Einklang zu bringen [275, 276].

aps Zugriffspfade werden eingesetzt, wann immer dies gewinnbringend möglich ist. Beispielsweise kann die Anfrage

```
select count(*) from R;
```

durch einen Indexscan effizient ausgewertet werden [145].

tmpidx Manchmal kann es sinnvoll sein, temporäre Zugriffspfade anzulegen.

optimpl Für algebraische Operatoren existieren im allgemeinen mehrere Implementierungen. Es sollte hier immer die für einen Operator im vorliegenden Fall billigste Lösung ausgewählt werden. Ebenfalls von Bedeutung ist die Darstellung des Zwischenergebnisses. Beispielsweise können Relationen explizit oder implizit dargestellt werden, wobei letztere Darstellung nur Zeiger auf Tupel oder Surrogate der Tupel enthält. Weitergedacht führt diese Technik zu den TID-Listen-basierten Operatoren.

setpipe Die Evaluation eines algebraischen Ausdrucks kann entweder mengenorientiert oder nebenläufig (pipelining) erfolgen. Letzteres erspart das Erzeugen von großen Zwischenergebnissen.

tmpplay Das temporäre Ändern eines Layouts eines Objektes kann durchaus sinnvoll sein, wenn die Kosten, die durch diese Änderung entstehen, durch den Gewinn der mehrmaligen Verwendung dieses Layouts mehr als kompensiert werden. Ein typisches Beispiel ist *Pointer-swizzling*.

matSubQ If a query is not unnested, then for every argument combination passed to the subquery, the result is materialized in order to avoid duplicate computation of the same subquery expression for the same argument combination [275, 276]. This technique is favorable for detachment [734, 808, 840]

AggrJoin Joins with non-equi join predicates based on \leq or $<$, can be processed more efficiently than by a cross product with a subsequent selection [164].

ClassHier Class hierarchies involve the computation of queries over a union of extents (if implemented that way). Pushing algebraic operations past unions allows often for more efficient plans [166].

AggrIDX Use an index to determine aggregate values like min/max/avg/count.

rid/tidsort When several tuples qualify during an index scan, the resulting TIDs can be sorted in order to guarantee sequential access to the base relation.

multIDX Perform operations like union and disjunction on the outcome of an index scan.

multIDXsplit If two ranges are queried within the same query ([1-10],[20-30]) consider multIDX or use a single scan through the index [1-30] with an additional qualification predicate.

multIDXor Queries with more conditions on indexed attributes can be evaluated by more complex combinations of index scans and tid-list/bitvector operations. ($A = 5$ and ($B = 3$ or $B = 4$)).

scanDirChange During multiple sequential scans of relation (e.g. for a block-wise nested loop join), the direction of the scan can be changed in order to reuse as much of the pages in the buffer as possible.

lock The optimizer should chose the correct locks to set on tables. For example, if a whole table is scanned, a table lock should be set.

expFunMat Expensive functions can be cached during query evaluation in order to avoid their multiple evaluation for the same arguments [355].

expFunFil Easier to evaluate predicates that are implied by more expensive predicates can serve as filters in order to avoid the evaluation of the expensive predicate on all tuples.

stop Stop evaluation after the first tuple qualifies. This is good for existential subqueries, universal subqueries (disqualify), semi-joins for distinct results and the like.

expensive projections 1. zum Schluss, da dort am wenigsten verschiedene Werte

2. durchschieben, falls cache fuer Funktionsergebnisse dadurch vermieden werden kann

OO-Kontext: problematisch: objekte muessen fuer funktionen/methoden als ganzes vorhanden sein. daher ist eine einfache strategie nicht moeglich.

distinct/sorting `select distinct a,b,c`
 `...`
 `order by a,b`

kann auch nach a,b,c sortiert werden. stoert gar nicht, vereinfacht aber die duplikateliminierung. nur ein sortieren notwendig.

index access • by key

- by key range
- by dashed key range (set of keys/key ranges)
- index anding/oring

alternative operator implementations e.g. join: nlj bnlj hj grace-hash
 hybrid-hash smj diag-join star-join

distpd Push-down or Pull-up distinct.

aggregate with distinct `select a, agg(distinct b)`

```
...
group by a
====>
sort on a,b
dup elim
group a,sum(b)
```

alternative: `aggr(distinct *)` is implemented such that it uses a hashtable to eliminate duplicates this is only good, if the number of groups is small and the number of distinct values in each group is small.

XXX - use keys, inclusion dependencies, fds etc. (all user specified and derived) (propagate keys over joins as fds), (for a function call: derived IU is functional dependent on arguments of the function call if function is deterministic) (keys can be represented as sets of IUs or as bitvectors (given numbering of IUs)) (numbering imprecise: bitvectors can be used as filters (like for signatures))

Appendix D

Useful Formulas

The following identities can be found in the book by Graham, Knuth, and Patashnik [308].

We use the following definition of binomial coefficients:

$$\binom{n}{k} = \begin{cases} \frac{n!}{k!(n-k)!} & \text{if } 0 \leq k \leq n \\ 0 & \text{else} \end{cases} \quad (\text{D.1})$$

We start with some simple identities.

$$\binom{n}{k} = \binom{n}{n-k} \quad (\text{D.2})$$

$$\binom{n}{k} = \frac{n}{k} \binom{n-1}{k-1} \quad (\text{D.3})$$

$$k \binom{n}{k} = n \binom{n-1}{k-1} \quad (\text{D.4})$$

$$(n-k) \binom{n}{k} = n \binom{n-1}{k} \quad (\text{D.5})$$

$$(n-k) \binom{n}{k} = n \binom{n-1}{n-k-1} \quad (\text{D.6})$$

$$\binom{n}{k} = \binom{n-1}{k} + \binom{n-1}{k-1} \quad (\text{D.7})$$

$$\binom{r}{m} \binom{m}{k} = \binom{r}{k} \binom{r-k}{m-k} \quad (\text{D.8})$$

The following identities are good for sums of binomial coefficients.

$$\sum_{k=0}^n \binom{n}{k} = 2^n \quad (\text{D.9})$$

$$\sum_{k=0}^n \binom{k}{m} = \binom{n+1}{m+1} \quad (\text{D.10})$$

$$\sum_{k=0}^n \binom{m+k}{k} = \binom{m+n+1}{m+1} = \binom{m+n+1}{n} \quad (\text{D.11})$$

$$\sum_{k=0}^n \binom{m-n+k}{k} = \binom{m+1}{n} \quad (\text{D.12})$$

From Identities D.2 and D.11 it follows that

$$\sum_{k=0}^m \binom{k+r}{r} = \binom{m+r+1}{r+1} \quad (\text{D.13})$$

For sums of products, we have

$$\sum_{k=0}^n \binom{r}{m+k} \binom{s}{n-k} = \binom{r+s}{m+n} \quad (\text{D.14})$$

$$\sum_{k=0}^n \binom{l-k}{m} \binom{q+k}{n} = \binom{l+q+1}{m+n+1} \quad (\text{D.15})$$

$$\sum_{k=0}^n \binom{l}{m+k} \binom{s}{n+k} = \binom{l+s}{l-m+n} \quad (\text{D.16})$$

Bibliography

- [1] K. Aberer and G. Fischer. Semantic query optimization for methods in object-oriented database systems. In *Proc. IEEE Conference on Data Engineering*, pages 70–79, 1995.
- [2] S. Abiteboul. On Views and XML. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages 1–9, 1999.
- [3] S. Abiteboul, C. Beeri, M. Gyssens, and D. Van Gucht. An introduction to the completeness of languages for complex objects and nested relations. In S. Abiteboul, P.C. Fischer, and H.-J. Schek, editors, *Nested Relations and Complex Objects in Databases*, pages 117–138. Lecture PAGESs in Computer Science 361, Springer, 1987.
- [4] S. Abiteboul and N. Bidoit. Non first normal form relations: An algebra allowing restructuring. *Journal of Computer Science and Systems*, 33(3):361, 1986.
- [5] S. Abiteboul, S. Cluet, V. Christophides, T. Milo, G. Moerkotte, and J. Simeon. Querying documents in object databases. *International Journal on Digital Libraries*, 1(1):5–19, April 1997.
- [6] S. Abiteboul and O. Duschka. Complexity of answering queries using materialized views. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages 254–263, 1998.
- [7] A. Abounaga, A. Alameldeen, and J. Naughton. Estimating the selectivity of XML path expressions for internet scale applications. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 591–600, 2001.
- [8] A. Abounaga and J. Naughton. Building XML statistics for the hidden web. In *Int. Conference on Information and Knowledge Management (CIKM)*, pages 358–365, 2003.
- [9] W. Abu-Sufah, D. J. Kuch, and D. H. Lawrie. On the performance enhancement of paging systems through program analysis and transformations. *IEEE Trans. on Computers*, C-50(5):341–355, 1981.
- [10] B. Adelberg, H. Garcia-Molina, and J. Widom. The STRIP rule system for efficiently maintaining derived data. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 147–158, 1997.

- [11] F. Afrati, M. Gergatsoulis, and T. Kavalieros. Answering queries using materialized views with disjunctions. In *Proc. Int. Conf. on Database Theory (ICDT)*, pages 435–452, 1999.
- [12] F. Afrati, C. Li, and J. Ullman. Generating efficient plans for queries using views. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 319–330, 2001.
- [13] F. Afrati and C. Papadimitriou. The parallel complexity of simple chain queries. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages 210–?, 1987.
- [14] D. Agrawal, A. Abbadi, A. Singh, and T. Yurek. Efficient view maintenance at data warehouses. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 417–427, 1997.
- [15] S. Agrawal, R. Agrawal, P. Deshpande, A. Gupta, J. Naughton, R. Ramakrishnan, and S. Sarawagi. On the computation of multidimensional aggregates. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 506–521, 1996.
- [16] A. Aho, Y. Sagiv, and J. Ullman. Efficient optimization of a class of relational expressions. *ACM Trans. on Database Systems*, 4(4):435–454, 1979.
- [17] A. Aho, Y. Sagiv, and J. Ullman. Equivalence of relational expressions. *SIAM J. on Computing*, 8(2):218–246, 1979.
- [18] A. V. Aho, C. Beeri, and J. D. Ullman. The theory of joins in relational databases. *ACM Trans. on Database Systems*, 4(3):297–314, 1979.
- [19] A. V. Aho, Y. Sagiv, and J. D. Ullman. Efficient optimization of a class of relational expression. *ACM Trans. on Database Systems*, 4(4):435–454, 1979.
- [20] A.V. Aho, Y. Sagiv, and J.D. Ullman. Equivalence among relational expressions. *SIAM Journal on Computing*, 8(2):218–246, 1979.
- [21] S. Al-Khalifa, H. Jagadish, N. Koudas, J. Patel, D. Srivastava, and Y. Wu. Structural joins: A primitive for efficient XML query pattern matching. In *Proc. IEEE Conference on Data Engineering*, pages 141–152, 2002.
- [22] J. Albert. Algebraic properties of bag data types. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 211–219, 1991.
- [23] F. E. Allen and J. Cocke. A catalogue of optimizing transformations. In R. Rustin, editor, *Design and optimization of compilers*, pages 1–30. Prentice Hall, 1971.
- [24] N. Alon, P. Gibbons, Y. Matias, and M. Szegedy. Tracking join and self-join sizes in limited storage. *J. Comput. Syst. Sci.*, 35(4):391–432, 2002.

- [25] J. Alsabbagh and V. Rahavan. Analysis of common subexpression exploitation models in multiple-query processing. In *Proc. IEEE Conference on Data Engineering*, pages 488–497, 1994.
- [26] P. Alsberg. Space and time savings through large database compression and dynamic restructuring. In *Proc IEEE 63,8*, Aug. 1975.
- [27] S. Amer-Yahia, S. Cho, L. Lakshmanan, and D. Srivastava. Efficient algorithms for minimizing tree pattern queries. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 497–508, 2001.
- [28] L. Amsaleg, M. Franklin, A. Tomasic, and T. Urhan. Scrambling query plans to cope with unexpected delay. In *4th Int. Conference on Parallel and Distributed Information Systems (PDIS)*, Palm Beach, Fl, 1996.
- [29] O. Anfinsen. A study of access path selection in DB2. Technical report, Norwegian Telecommunication Administration and University of Oslo, Norway, Oct. 1989.
- [30] G. Antoshenkov. Dynamic query optimization in RDB/VMS. In *Proc. IEEE Conference on Data Engineering*, pages 538–547, Vienna, Apr. 1993.
- [31] G. Antoshenkov. Query processing in DEC Rdb: Major issues and future challenges. *IEEE Data Engineering Bulletin*, 16:42–52, Dec. 1993.
- [32] G. Antoshenkov. Dynamic optimization of index scans restricted by booleans. In *Proc. IEEE Conference on Data Engineering*, pages 430–440, 1996.
- [33] P.M.G. Apers, A.R. Hevner, and S.B. Yao. Optimization algorithms for distributed queries. *IEEE Trans. on Software Eng.*, 9(1):57–68, 1983.
- [34] P.M.G. Apers, A.R. Hevner, and S.B. Yao. Optimization algorithms for distributed queries. *IEEE Trans. on Software Eng.*, 9(1):57–68, 1983.
- [35] R. Ashenurst. Acm forum. *Communications of the ACM*, 20(8):609–612, 1977.
- [36] M. M. Astrahan and D. D. Chamberlin. Implementation of a structured English query language. *Communications of the ACM*, 18(10):580–588, 1975.
- [37] M.M. Astrahan, M.W. Blasgen, D.D. Chamberlin, K.P. Eswaran, J.N. Gray, P.P. Griffiths, W.F. King, R.A. Lorie, P.R. Mc Jones, J.W. Mehl, G.R. Putzolu, I.L. Traiger, B.W. Wade, and V. Watson. System R: relational approach to database management. *ACM Transactions on Database Systems*, 1(2):97–137, June 1976.
- [38] R. Avnur and J. Hellerstein. Eddies: Continuously adaptive query optimization. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, 2000.

- [39] B. Babcock and S. Chaudhuri. Towards a robust query optimizer: A principled and practical approach. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 119–130, 2005.
- [40] S. Babu, P. Bizarro, and D. DeWitt. Proactive re-optimization. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 107–118, 2005.
- [41] L. Baekgaard and L. Mark. Incremental computation of nested relational query expressions. *ACM Trans. on Database Systems*, 20(2):111–148, 1995.
- [42] T. Baer. Iperfex: A hardware performance monitor for Linux/IA32 systems. perform internet search for this or similar tools.
- [43] A. Balmin, F. Özcan, K. Beyer, R. Cochrane, and H. Pirahesh. A framework for using materialized XPath views in XML query processing. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 60–71, 2004.
- [44] F. Bancilhon and K. Koshafian. A calculus for complex objects. In *ACM Symp. on Principles of Database Systems*, pages 53–59, 1986.
- [45] J. Banerjee, W. Kim, and K.-C. Kim. Queries in object-oriented databases. MCC Technical Report DB-188-87, MCC, Austin, TX 78759, June 1987.
- [46] J. Banerjee, W. Kim, and K.-C. Kim. Queries in object-oriented databases. In *Proc. IEEE Conference on Data Engineering*, pages 31–38, 1988.
- [47] E. Baralis, S. Paraboschi, and E. Teniente. Materialized views selection in a multidimensional database. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 156–165, 1997.
- [48] M. Bassiouni. Data compression in scientific and statistical databases. *IEEE Trans. on Software Eng.*, 11(10):1047–1058, 1985.
- [49] D. Batory. On searching transposed files. *ACM Trans. on Database Systems*, 4(4):531–544, 1979.
- [50] D. Batory. Extensible cost models and query optimization in Genesis. *IEEE Database Engineering*, 9(4), Nov 1986.
- [51] D. S. Batory. Modeling the storage architecture of commercial database systems. *ACM Trans. on Database Systems*, 10(4):463–528, Dec. 1985.
- [52] D. S. Batory. A molecular database systems technology. Tech. Report TR-87-23, University of Austin, 1987.
- [53] D. S. Batory. Building blocks of database management systems. Technical Report TR-87-23, University of Texas, Austin, TX, Feb. 1988.
- [54] D. S. Batory. Concepts for a database system compiler. In *Proc. of the 17th ACM SIGMOD*, pages 184–192, 1988.

- [55] D. S. Batory. On the reusability of query optimization algorithms. *Information Sciences*, 49:177–202, 1989.
- [56] D. S. Batory and C. Gotlieb. A unifying model of physical databases. *ACM Trans. on Database Systems*, 7(4):509–539, Dec. 1982.
- [57] D. S. Batory, T. Y. Leung, and T. E. Wise. Implementation concepts for an extensible data model and data language. *ACM Trans. on Database Systems*, 13(3):231–262, Sep 1988.
- [58] L. Baxter. *The Complexity of Unification*. PhD thesis, University of Waterloo, 1976.
- [59] L. Becker and R. H. Güting. Rule-based optimization and query processing in an extensible geometric database system. *ACM Trans. on Database Systems (to appear)*, 1991.
- [60] L. Becker and R. H. Güting. Rule-based optimization and query processing in an extensible geometric database system. *ACM Trans. on Database Systems*, 17(2):247–303, June 1992.
- [61] C. Beeri and Y. Kornatzky. Algebraic optimization of object-oriented query languages. In *Proc. Int. Conf. on Database Theory (ICDT)*, pages 72–88, 1990.
- [62] C. Beeri and Y. Kornatzky. Algebraic optimization of object-oriented query languages. *Theoretical Computer Science*, 116(1):59–94, 1993.
- [63] C. Beeri and Y. Tzaban. SAL: An algebra for semistructured data and XML. In *ACM SIGMOD Workshop on the Web and Databases (WebDB)*, 1999.
- [64] L. A. Belady. A study of replacement algorithms for virtual storage computers. *IBM Systems Journal*, 5(2):78–101, 1966.
- [65] D. Beneventano, S. Bergamaschi, and C. Sartori. Description logic for semantic query optimization in object-oriented database systems. *ACM Trans. on Database Systems*, 28(1):1–50, 2003.
- [66] K. Bennett, M. Ferris, and Y. Ioannidis. A genetic algorithm for database query optimization. Technical Report Tech. Report 1004, University of Wisconsin, 1990.
- [67] K. Bennett, M. Ferris, and Y. Ioannidis. A genetic algorithm for database query optimization. In *Proc. 4th Int. Conf. on Genetic Algorithms*, pages 400–407, 1991.
- [68] A. Bernasconi and B. Codenetti. Measures of boolean function complexity based on harmonic analysis. In M. Bonuccelli, P. Crescenzi, and R. Petreschi, editors, *Algorithms and Complexity (2nd. Italian Conference)*, pages 63–72, Rome, Feb. 1994. Springer (Lecture Notes in Computer Science 778).

- [69] P. Bernstein, E. Wong, C. Reeve, and J. Rothnie. Query processing in a system for distributed databases (sdd-1). *ACM Trans. on Database Systems*, 6(4):603–625, 1981.
- [70] P. A. Bernstein and D. M. W. Chiu. Using semi-joins to solve relational queries. *Journal of the ACM*, 28(1):25–40, 1981.
- [71] P. A. Bernstein and N. Goodman. The power of inequality semijoin. *Information Systems*, 6(4):255–265, 1981.
- [72] P. A. Bernstein and N. Goodman. The power of natural semijoin. *SIAM J. Comp.*, 10(4):751–771, 1981.
- [73] E. Bertino and D. Musto. Query optimization by using knowledge about data semantics. *Data & Knowledge Engineering*, 9(2):121–155, 1992.
- [74] E. Bertino, M. Negri, G. Pelagatti, and L. Sbattella. Object-oriented query languages: The notion and the issues. *IEEE Trans. on Knowledge and Data Eng.*, 4(3):223–237, June 1992.
- [75] G. Bhargava, P. Goel, and B. Iyer. Hypergraph based reorderings of outer join queries with complex predicates. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 304–315, 1995.
- [76] G. Bhargava, P. Goel, and B. Iyer. Efficient processing of outer joins and aggregate functions. In *Proc. IEEE Conference on Data Engineering*, pages 441–449, 1996.
- [77] A. Biliris. An efficient database storage structure for large dynamic objects. In *Proc. IEEE Conference on Data Engineering*, pages 301–308, 1992.
- [78] D. Bitton and D. DeWitt. Duplicate record elimination in large data files. *ACM Trans. on Database Systems*, 8(2):255–265, 1983.
- [79] J. Blakeley and N. Martin. Join index, materialized view, and hybrid hash-join: a performance analysis. In *Proc. IEEE Conference on Data Engineering*, pages 256–236, 1990.
- [80] J. Blakeley, W. McKenna, and G. Graefe. Experiences building the Open OODB query optimizer. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 287–295, 1993.
- [81] J. A. Blakeley, P. A. Larson, and F. W. Tompa. Efficiently updating materialized views. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 61–71, Washington, D.C., 1986.
- [82] P. Bodorik and J.S. Riordon. Distributed query processing optimization objectives. In *Proc. IEEE Conference on Data Engineering*, pages 320–329, 1988.

- [83] T. Böhme and E. Rahm. Xmach-1: A benchmark for XML data management. In *BTW*, pages 264–273, 2001.
- [84] A. Bolour. Optimal retrieval for small range queries. *SIAM J. of Comput.*, 10(4):721–741, 1981.
- [85] P. Bonatti. On the decidability of containment of recursive datalog queries - preliminary report. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages 297–306, 2004.
- [86] P. Boncz, A. Wilschut, and M. Kersten. Flattening an object algebra to provide performance. In *Proc. IEEE Conference on Data Engineering*, pages 568–577, 1998.
- [87] B. Bouchou, M. Halfeld, and F. Alves. Updates and incremental validation of XML documents. In *Int. Workshop on Database Programming Languages*, pages 216–232, 2003.
- [88] M. Brantner, S. Helmer, C.-C. Kanne, and G. Moerkotte. Full-fledged algebraic XPath processing in Natix. In *Proc. IEEE Conference on Data Engineering*, pages 705–716, 2005.
- [89] K. Bratbergsengen and K. Norvag. Improved and optimized partitioning techniques in database query processing. In *Advances in Databases, 15th British National Conference on Databases*, pages 69–83, 1997.
- [90] Y. Breitbart and A. Reiter. Algorithms for fast evaluation of boolean expressions. *Acta Informatica*, 4:107–116, 1975.
- [91] S. Bressan, M. Lee, Y. Li, Z. Lacroix, and U. Nambiar. The XOO7 XML Management System Benchmark. Technical Report TR21/00, National University of Singapore, 2001.
- [92] K. Brown, M. Carey, and M. Livny. Goal-oriented buffer management revisited. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 353–364, Montreal, Canada, Jun 1996.
- [93] F. Bry. Logical rewritings for improving the evaluation of quantified queries. In *2nd. Symp. on Mathematical Fundamentals of Database Systems*, pages 100–116, June 1989, Visegrad, Hungary, 1989.
- [94] F. Bry. Towards an efficient evaluation of general queries: Quantifiers and disjunction processing revisited. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 193–204, 1989.
- [95] P. Buneman, S. Davidson, W. Fan, C. Hara, and W. Tan. Keys for XML. In *WWW Conference*, pages 201–210, 2001.
- [96] P. Buneman, L. Libkin, D. Suciu, V. Tannen, and L. Wong. Comprehension syntax. *SIGMOD Record*, 23(1):87–96, 1994.

- [97] L. Cabibbo and R. Torlone. A framework for the investigation of aggregate functions in database queries. In *Proc. Int. Conf. on Database Theory (ICDT)*, pages 383–397, 1999.
- [98] J.-Y. Cai, V. Chakaravarthy, R. Kaushik, and J. Naughton. On the complexity of join predicates. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages 207–214, 2001.
- [99] D. Calvanese, G. DeGiacomo, M. Lenzerini, and M. Vardi. View-based query answering and query containment over semistructured data. In *Int. Workshop on Database Programming Languages*, pages 40–61, 2001.
- [100] B. Cao and A. Badia. A nested relational approach to processing sql subqueries. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 191–202, 2005.
- [101] L. Cardelli and P. Wegner. On understanding types, data abstraction, and polymorphism. *JOURNAL of the ACM*, Computing Surveys:471–522, 1985.
- [102] A. F. Cardenas. Analysis and performance of inverted data base structures. *Communications of the ACM*, 18(5):253–263, 1975.
- [103] M. Carey, D. DeWitt, J. Richardson, and E. Shikita. Object and file management in the EXODUS extensible database system. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 91–100, 1986.
- [104] M. Carey and D. Kossmann. On saying “enough already!” in SQL. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 219–230, 1997.
- [105] M. Carey and D. Kossmann. Processing top N and bottom N queries. *IEEE Data Engineering Bulletin*, 20(3):12–19, 1997.
- [106] M. Carey and D. Kossmann. Reducing the braking distance of an SQL query engine. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 158–169, 1998.
- [107] J. Carlis. HAS: A relational algebra operator, or divide is not to conquer. In *Proc. IEEE Conference on Data Engineering*, pages 254–261, 1986.
- [108] L. Carlitz, D. Roselle, and R. Scoville. Some remarks on ballot-type sequences of positive integers. *Journal of Combinatorial Theory*, 11:258–271, 1971.
- [109] C. R. Carlson and R. S. Kaplan. A generalized access path model and its application to a relational database system. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 143–154, 1976.
- [110] S. Ceri and G. Gottlob. Translating SQL into relational algebra: Optimization, semantics and equivalence of SQL queries. *IEEE Trans. on Software Eng.*, 11(4):324–345, Apr 1985.

- [111] S. Ceri and G. Pelagatti. Correctness of query execution strategies in distributed databases. *ACM Trans. on Database Systems*, 8(4):577–607, Dec. 1983.
- [112] S. Ceri and G. Pelagatti. *Distributed Databases: Principles and Systems*. McGraw-Hill, 1985.
- [113] S. Chakravarthy. Devide and conquer: a basis for augmenting a conventional query optimizer with multiple query processing capabilities. In *Proc. IEEE Conference on Data Engineering*, 1991.
- [114] U. S. Chakravarthy, J. Grant, and J. Minker. Logic-based approach to semantic query optimization. *ACM Trans. on Database Systems*, 15(2):162–207, 1990.
- [115] U. S. Chakravarthy and J. Minker. Multiple query processing in deductive databases using query graphs. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages ?–?, 1986.
- [116] D. Chamberlin, J. Robie, and D. Florescu. Quilt: An XML query language for heterogeneous data sources. In *ACM SIGMOD Workshop on the Web and Databases (WebDB)*, 2000.
- [117] C. Chan and B. Ooi. Efficient scheduling of page accesses in index-based join processing. *IEEE Trans. on Knowledge and Data Eng.*, 9(6):1005–1011, 1997.
- [118] A. Chandra and P. Merlin. Optimal implementation of conjunctive queries in relational data bases. In *Proc. ACM SIGACT Symp. on the Theory of Computing*, pages 77–90, 1977.
- [119] S. Chatterji, S. Evani, S. Ganguly, and M. Yemmanuru. On the complexity of approximate query optimization. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages 282–292, 2002.
- [120] D. Chatziantoniou, M. Akinde, T. Johnson, and S. Kim. The MD-Join: An Operator for Complex OLAP. In *Proc. IEEE Conference on Data Engineering*, pages 524–533, 2001.
- [121] D. Chatziantoniou and K. Ross. Querying multiple features in relational databases. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 295–306, 1996.
- [122] D. Chatziantoniou and K. Ross. Groupwise processing of relational queries. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 476–485, 1997.
- [123] S. Chaudhuri, P. Ganesan, and S. Sarawagi. Factorizing complex predicates in queries to exploit indexes. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 361–372, 2003.

- [124] S. Chaudhuri and L. Gravano. Evaluating top-k selection queries. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 397–410, 1999.
- [125] S. Chaudhuri, R. Krishnamurthy, S. Potamianos, and K. Shim. Optimizing queries with materialized views. In *Proc. IEEE Conference on Data Engineering*, pages 190–200, 1995.
- [126] S. Chaudhuri, R. Krishnamurthy, S. Potamianos, and K. Shim. *Optimizing Queries with Materialized Views*, pages 77–92. MIT Press, 1999.
- [127] S. Chaudhuri, V. Narasayya, and R. Ramamurthy. Estimating progress of long running sql queries. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 803–814, 2004.
- [128] S. Chaudhuri and K. Shim. Query optimization in the presence of foreign functions. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 529–542, 1993.
- [129] S. Chaudhuri and K. Shim. Including group-by in query optimization. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 354–366, 1994.
- [130] S. Chaudhuri and K. Shim. The promise of early aggregation. Technical report, HP Lab, 1994. Never Appeared.
- [131] S. Chaudhuri and K. Shim. An overview of cost-based optimization of queries with aggregates. *IEEE Data Engineering Bulletin*, 18(3):3–9, Sept 1995.
- [132] S. Chaudhuri and K. Shim. Optimization of queries with user-defined predicates. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 87–98, 1996.
- [133] S. Chaudhuri and K. Shim. Optimizing queries with aggregate views. In *Proc. of the Int. Conf. on Extending Database Technology (EDBT)*, pages 167–182, 1996.
- [134] S. Chaudhuri and M. Vardi. On the equivalence of recursive and nonrecursive datalog programs. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages 55–66, 1992.
- [135] S. Chaudhuri and M. Vardi. Optimization of real conjunctive queries. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages 59–70, 1993.
- [136] S. Chaudhuri and M. Vardi. Optimization of real conjunctive queries. Technical Report HPL-93-26, HP Software Technology Laboratory, 1993.
- [137] S. Chaudhuri and M. Vardi. On the complexity of equivalence between recursive and nonrecursive datalog programs. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages 55–66, 1994.

- [138] P. Cheeseman, B. Kanefsky, and W. Taylor. Where the *really* hard problems are. In *Int. Joint Conf. on Artificial Intelligence (IJCAI)*, pages 331–337, 1991.
- [139] C. Chekuri and A. Rajaraman. Conjunctive query containment revisited. In *Proc. Int. Conf. on Database Theory (ICDT)*, pages 56–70, 1997.
- [140] C. Chen and N. Roussopoulos. The implementation and performance evaluation of the ADMS query optimizer: Integrating query result caching and matching. In *Proc. of the Int. Conf. on Extending Database Technology (EDBT)*, pages 323–336, 1994.
- [141] Z. Chen, J. Gehrke, and F. Korn. Query optimization in compressed database systems. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 271–282, 2001.
- [142] Z. Chen, H. V. Jagadish, F. Korn, N. Koudas, S. Muthukrishnan, R. Ng, and D. Srivastava. Counting twig matches in a tree. In *Proc. IEEE Conference on Data Engineering*, pages 595–604, 2001.
- [143] Z. Chen and V. Narasayya. Efficient computation of multiple group by queries. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 263–274, 2005.
- [144] J. Cheng, D. Haderle, R. Hedges, B. Iyer, T. Messenger, C. Mohan, and Y. Wang. An efficient hybrid join algorithm: A DB2 prototype. In *Proc. IEEE Conference on Data Engineering*, pages 171–180, 1991.
- [145] J. Cheng, C. Loosley, A. Shibamiya, and P. Worthington. IBM DB2 Performance: design, implementation, and tuning. *IBM Sys. J.*, 23(2):189–210, 1984.
- [146] Q. Cheng, J. Gryz, F. Koo, T. Y. Cliff Leung, L. Liu, X. Quian, and B. Schiefer. Implementation of two semantic query optimization techniques in DB2 universal database. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 687–698, 1999.
- [147] M. Cherniack and S. Zdonik. Rule languages and internal algebras for rule-based optimizers. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 401–412, 1996.
- [148] T.-Y. Cheung. Estimating block accesses and number of records in file management. *Communications of the ACM*, 25(7):484–487, 1982.
- [149] D. M. Chiu and Y. C. Ho. A methodology for interpreting tree queries into optimal semi-join expressions. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 169–178, 1980.
- [150] H.-T. Chou and D. DeWitt. An evaluation of buffer management strategies for relational database systems. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 127–141, 1985.

- [151] S. Christodoulakis. Estimating block transfers and join sizes. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 40–54, 1983.
- [152] S. Christodoulakis. Analysis of retrieval performance for records and objects using optical disk technology. *ACM Trans. on Database Systems*, 12(2):137–169, 1987.
- [153] V. Christophides, S. Cluet, and G. Moerkotte. Evaluating queries with generalized path expressions. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 413–422, 1996.
- [154] C. Clarke, G. Cormack, and F. Burkowski. An algebra for structured text search and a framework for its implementation. *The Computer Journal*, 38(1):43–56, 1995.
- [155] J. Claussen, A. Kemper, and D. Kossmann. Order-preserving hash joins: Sorting (almost) for free. Technical Report MIP-9810, University of Passau, 1998.
- [156] J. Claussen, A. Kemper, G. Moerkotte, and K. Peithner. Optimizing queries with universal quantification in object-oriented and object-relational databases. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 286–295, 1997.
- [157] J. Claussen, A. Kemper, G. Moerkotte, and K. Peithner. Optimizing queries with universal quantification in object-oriented and object-relational databases. Technical Report MIP-9706, University of Passau, Fak. f. Mathematik u. Informatik, Mar 1997.
- [158] J. Claussen, A. Kemper, G. Moerkotte, K. Peithner, and M. Steinbrunn. Optimization and evaluation of disjunctive queries. *IEEE Trans. on Knowledge and Data Eng.*, 12(2):238–260, 2000.
- [159] S. Cluet and C. Delobel. A general framework for the optimization of object-oriented queries. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 383–392, 1992.
- [160] S. Cluet, C. Delobel, C. Lecluse, and P. Richard. Reloop, an algebra based query language for an object-oriented database system. In *Proc. Int. Conf. on Deductive and Object-Oriented Databases (DOOD)*, 1989.
- [161] S. Cluet and G. Moerkotte. Nested queries in object bases. In *Proc. Int. Workshop on Database Programming Languages*, pages 226–242, 1993.
- [162] S. Cluet and G. Moerkotte. Classification and optimization of nested queries in object bases. Technical Report 95-6, RWTH Aachen, 1995.
- [163] S. Cluet and G. Moerkotte. Efficient evaluation of aggregates on bulk types. Technical Report 95-5, RWTH-Aachen, 1995.

- [164] S. Cluet and G. Moerkotte. Efficient evaluation of aggregates on bulk types. In *Proc. Int. Workshop on Database Programming Languages*, 1995.
- [165] S. Cluet and G. Moerkotte. On the complexity of generating optimal left-deep processing trees with cross products. In *Proc. Int. Conf. on Database Theory (ICDT)*, pages 54–67, 1995.
- [166] S. Cluet and G. Moerkotte. Query optimization techniques exploiting class hierarchies. Technical Report 95-7, RWTH-Aachen, 1995.
- [167] E. Codd. *Database Systems - Courant Computer Science Symposium*. Prentice Hall, 1972.
- [168] E. F. Codd. A database sublanguage founded on the relational calculus. In *Proc. ACM-SIGFIDET Workshop, Datadescription, Access, and Control*, pages 35–68, San Diego, Calif., 1971. ACM.
- [169] E. F. Codd. Relational completeness of data base sublanguages. In *Courant Computer Science Symposia No. 6: Data Base Systems*, pages 67–101, New York, 1972. Prentice Hall.
- [170] E. F. Codd. Extending the relational database model to capture more meaning. *ACM Trans. on Database Systems*, 4(4):397–434, Dec. 1979.
- [171] S. Cohen, W. Nutt, and Y. Sagiv. Equivalences among aggregate queries with negation. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages 215–226, 2001.
- [172] S. Cohen, W. Nutt, and Y. Sagiv. Containment of aggregate queries (extended version). Technical report, Hebrew University of Jerusalem, 2002. available at www.cs.huji.ac.il/~sarina/papers/agg-containment-long.ps.
- [173] S. Cohen, W. Nutt, and Y. Sagiv. Containment of aggregate queries. In *Proc. Int. Conf. on Database Theory (ICDT)*, pages 111–125, 2003.
- [174] L. Colby. A recursive algebra and query optimization for nested relational algebra. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 273–283, 1989.
- [175] L. Colby, A. Kawaguchi, D. Lieuwen, I. Mumick, and K. Ross. Supporting multiple view maintenance policies. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 405–416, 1997.
- [176] G. Copeland and S. Khoshafian. A decomposition storage model. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 268–279, Austin, TX, 1985.
- [177] G. Cormack. Data compression on a database system. *Communications of the ACM*, 28(12):1336–1342, 1985.

- [178] T. Cormen, C. Leiserson, and R. Rivest. *Introduction to Algorithms*. MIT Press, 1990.
- [179] T. Cormen, C. Leiserson, R. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, 2001. 2nd Edition.
- [180] D. Cornell and P. Yu. Integration of buffer management and query optimization in relational database environments. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 247–255, 1989.
- [181] J. Crawford and L. Auton. Experimental results on the crossover point in satisfiability problems. In *Proc. National Conference on Artificial Intelligence*, pages 21–27, 1993.
- [182] K. Culik, T. Ottmann, and D. Wood. Dense multiway trees. *ACM Trans. on Database Systems*, 6(3):486–512, 1981.
- [183] M. Dadashzadeh. An improved division operator for relational algebra. *Information Systems*, 14(5):431–437, 1989.
- [184] D. Daniels. Query compilation in a distributed database system. Technical Report RJ 3432, IBM Research Laboratory, San Jose, CA, 1982.
- [185] D. Das and D. Batory. Praire: A rule specification framework for query optimizers. In *Proc. IEEE Conference on Data Engineering*, pages 201–210, 1995.
- [186] C. J. Date. The outer join. In *Proc. of the Int. Conf. on Databases*, Cambridge, England, 1983.
- [187] U. Dayal. Processing queries with quantifiers: A horticultural approach. In *ACM Symp. on Principles of Database Systems*, pages 125–136, 1983.
- [188] U. Dayal. Of nests and trees: A unified approach to processing queries that contain nested subqueries, aggregates, and quantifiers. In *VLDB*, pages 197–208, 1987.
- [189] U. Dayal, N. Goodman, and R.H. Katz. An extended relational algebra with control over duplicate elimination. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages 117–123, 1982.
- [190] U. Dayal, F. Manola, A. Buchman, U. Chakravarthy, D. Goldhirsch, S. Heiler, J. Orenstein, and A. Rosenthal. Simplifying complex object: The PROBE approach to modelling and querying them. In *H.J. Schek and G. Schlageter (eds.) Proc. BTW*, pages 17–37, 1987.
- [191] U. Dayal and J. Smith. PROBE: A knowledge-oriented database management system. In *Proc. Islamorada Workshop on Large Scale Knowledge Base and Reasoning Systems*, 1985.

- [192] G. de Balbine. Note on random permutations. *Mathematics of Computation*, 21:710–712, 1967.
- [193] D. DeHaan, P.-A. Larson, and J. Zhou. Stacked index views in Microsoft SQL server. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 179–190, 2005.
- [194] R. Demolombe. Estimation of the number of tuples satisfying a query expressed in predicate calculus language. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 55–63, 1980.
- [195] P. Denning. Effects of scheduling on file memory operations. In *Proc. AFIPS*, pages 9–21, 1967.
- [196] N. Derrett and M.-C. Shan. Rule-based query optimization in IRIS. Technical report, Hewlett-Packard Laboratories, 1501 Page Mill Road, Palo Alto, CA94303, 1990.
- [197] B. Desai. Performance of a composite attribute and join index. *IEEE Trans. on Software Eng.*, 15(2):142–152, Feb. 1989.
- [198] A. Deshpande, M. Garofalakis, and R. Rastogi. Independence is good: Dependency-based histogram synopses for high-dimensional data. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 199–210, 2001.
- [199] A. Deutsch, M. Fernandez, D. Florescu, A. Levy, D. Maier, and D. Suciu. Querying XML data. *IEEE Data Engineering Bulletin*, 22(3):10–18, 1999.
- [200] A. Deutsch, B. Ludäscher, and A. Nash. Rewriting queries using views with access patterns under integrity constraints. In *Proc. Int. Conf. on Database Theory (ICDT)*, pages 352–367, 2005.
- [201] A. Deutsch and V. Tannen. Containment and integrity constraints for XPath. In *Int. Workshop on Knowledge Representation meets Databases (KRDB)*, 2001.
- [202] A. Deutsch and V. Tannen. Optimization properties for classes of conjunctive regular path queries. In *Int. Workshop on Database Programming Languages*, pages 21–39, 2001.
- [203] A. Deutsch and V. Tannen. Reformulation of XML queries and constraints. In *Proc. Int. Conf. on Database Theory (ICDT)*, pages 225–241, 2003.
- [204] D. DeWitt, R. Katz, F. Olken, L. Shapiro, M. Stonebraker, and D. Wood. Implementation techniques for main memory database systems. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 151–1??, 1984.
- [205] D. DeWitt, J. Naughton, and D. Schneider. An evaluation of non-equijoin algorithms. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 443–452, 1991.

- [206] Y. Diao, M. Altinel, M. Franklin, H. Zhang, and P. Fischer. Path sharing and predicate evaluation for high-performance xml filtering. *ACM Trans. on Database Systems*, 28(4):367–516, 2003.
- [207] G. Diehr and A. Saharia. Estimating block accesses in database organizations. *IEEE Trans. on Knowledge and Data Engineering*, 6(3):497–499, 1994.
- [208] P. Dietz. Optimal algorithms for list indexing and subset ranking. In *Workshop on Algorithms and Data Structures (LNCS 382)*, pages 39–46, 1989.
- [209] Z. Dimitrijevic, R. Rangaswami, E. Chang, D. Watson, and A. Acharya. Diskbench: User-level disk feature extraction tool. Technical report, University of California, Santa Barbara, 2004.
- [210] J. Dongarra, K. London, S. Moore, P. Mucci, and D. Terpstra. Using PAPI for hardware performance monitoring on Linux systems. perform internet search for this or similar tools.
- [211] D. Donjerkovic and R. Ramakrishnan. Probabilistic optimization of Top n queries. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 411–422, 1999.
- [212] J. Donovan. Database system approach to management of decision support. *ACM Trans. on Database Systems*, 1(4):344–368, 1976.
- [213] M. Drmota, D. Gardy, and B. Gittenberger. A unified presentation of some urn models. *Algorithmica*, 29:120–147, 2001.
- [214] R. Durstenfeld. Algorithm 235: Random permutation. *Communications of the ACM*, 7(7):420, 1964.
- [215] O. Duschka. *Query Planning and Optimization in Information Integration*. PhD thesis, Stanford University, 1997.
- [216] O. Duschka and M. Genesereth. Answering recursive queries using views. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages 109–116, 1997.
- [217] O. Duschka and M. Genesereth. Query planning with disjunctive sources. In *AAAI Workshop on AI and Information Integration*, 1998.
- [218] O. Duschka and M. Gensereth. Answering recursive queries using views. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages 109–116, 1997.
- [219] O. Duschka and A. Levy. Recursive plans for information gathering. In *Int. Joint Conf. on Artificial Intelligence (IJCAI)*, 1997.
- [220] W. Effelsberg and T. Härder. Principles of database buffer management. *ACM Trans. on Database Systems*, 9(4):560–595, 1984.

- [221] S. Eggers, F. Olken, and A. Shoshani. A compression technique for large statistical data bases. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 424–434, 1981.
- [222] S. Eggers and A. Shoshani. Efficient access of compressed data. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 205–211, 1980.
- [223] J. F. Egler. A procedure for converting logic table conditions into an efficient sequence of test instructions. *Communications of the ACM*, 6(9):510–514, 1963.
- [224] M. Eisner and D. Severance. Mathematical techniques for efficient record segmentation in large shared databases. *Journal of the ACM*, 23(4):619–635, 1976.
- [225] R. Elmasri and S. Navathe. *Fundamentals of Database Systems*. Addison-Wesley, 2000. 3rd Edition.
- [226] R. Epstein. Techniques for processing of aggregates in relational database systems. ERL/UCB Memo M79/8, University of California, Berkeley, 1979.
- [227] G. Lohman et al. Optimization of nested queries in a distributed relational database. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, 1984.
- [228] N. Roussopoulos et al. The maryland ADMS project: Views R Us. *IEEE Data Engineering Bulletin*, 18(2), 1995.
- [229] P. Schwarz et al. Extensibility in the starburst database system. In *Proc. Int. Workshop on Object-Oriented Database Systems*, 1986.
- [230] R. Fagin. Combining fuzzy information from multiple systems. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages 216–226, 1996.
- [231] R. Fagin, A. Lotem, and M. Naor. Optimal aggregation algorithms for middleware. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages ?–?, 2001.
- [232] C. Farr'e, E. Teniente, and T. Urp'i. Query containment checking as a view updating problem. In *Inf. Conf. on Database and Expert Systems Applications (DEXA)*, pages 310–321, 1998.
- [233] C. Farr'e, E. Teniente, and T. Urp'i. The constructive method for query containment checking. In *Inf. Conf. on Database and Expert Systems Applications (DEXA)*, pages 583–593, 1999.
- [234] C. Farr'e, E. Teniente, and T. Urp'i. Query containment with negated IDB predicates. In *ABDIS*, pages 583–593, 2003.
- [235] C. Farr'e, E. Teniente, and T. Urp'i. Checking query containment with the cqc method. *Data & Knowledge Engineering*, 53:163–223, 2005.

- [236] L. Fegaras. Optimizing large OODB queries. In *Proc. Int. Conf. on Deductive and Object-Oriented Databases (DOOD)*, pages 421–422, 1997.
- [237] L. Fegaras. A new heuristic for optimizing large queries. In *DEXA*, pages 726–735, 1998.
- [238] L. Fegaras and D. Maier. Towards an effective calculus for object query languages. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 47–58, 1995.
- [239] L. Fegaras and D. Maier. Optimizing object queries using an effective calculus. *ACM Trans. on Database Systems*, 25(4):457–516, 2000.
- [240] M. Fernandez, A. Malhotra, J. Marsh, M. Nagy, and N. Walsh. Xquery 1.0: An xml query language. Technical report, W3C, 2003. W3C Working Draft.
- [241] T. Fiebig and G. Moerkotte. Evaluating Queries on Structure with eXtended Access Support Relations. In *WebDB 2000*, 2000.
- [242] T. Fiebig and G. Moerkotte. Algebraic XML construction in Natix. In *Proc. Int. Conf. on Web Information Systems Engineering (WISE)*, pages 212–221, 2001.
- [243] T. Fiebig and G. Moerkotte. Algebraic XML construction and its optimization in Natix. *World Wide Web Journal*, 4(3):167–187, 2002.
- [244] S. Finkelstein. Common expression analysis in database applications. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 235–245, 1982.
- [245] D. Florescu. *Espaces de Recherche pour l’Optimisation de Requêtes Objet (Search Spaces for Query Optimization)*. PhD thesis, Université de Paris VI, 1996. in French.
- [246] D. Florescu, A. Levy, and D. Suciu. Query containment for conjunctive queries with regular expressions. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages 139–148, 1998.
- [247] P. Fortier. *SQL-3, Implementing the SQL Foundation Standard*. McGraw Hill, 1999.
- [248] F. Fotouhi and S. Pramanik. Optimal secondary storage access sequence for performing relational join. *IEEE Trans. on Knowledge and Data Eng.*, 1(3):318–328, 1989.
- [249] F. Frasincar, G.-J. Houben, and C. Pau. XAL: An algebra for XML query optimization. In *Australasian Database Conference (ADC)*, 2002.
- [250] J. Freire, J. Haritsa, M. Ramanath, P. Roy, and J. Simeon. StatiX: making XML count. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 181–191, 2002.

- [251] J. C. Freytag. A rule-based view of query optimization. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 173–180, 1987.
- [252] J. C. Freytag and N. Goodman. Translating aggregate queries into iterative programs. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages ?–?, 1986.
- [253] J. C. Freytag and N. Goodman. On the translation of relational queries into iterative programs. *ACM Trans. on Database Systems*, 14(1):1–27, 1989.
- [254] C. Galindo-Legaria. Outerjoins as disjunctions. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 348–358, 1994.
- [255] C. Galindo-Legaria. Outerjoins as disjunctions. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 348–538, 1994.
- [256] C. Galindo-Legaria and M. Joshi. Orthogonal optimization of subqueries and aggregation. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 571–581, 2001.
- [257] C. Galindo-Legaria, A. Pellenkoft, and M. Kersten. Cost distribution of search spaces in query optimization. Technical Report CS-R9432, CWI, Amsterdam, NL, 1994.
- [258] C. Galindo-Legaria, A. Pellenkoft, and M. Kersten. Fact, randomized join-order selection — why use transformations? In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 85–95, 1994.
- [259] C. Galindo-Legaria, A. Pellenkoft, and M. Kersten. Fast, randomized join-order selection — why use transformations? Technical Report CS-R-9416, CWI, Amsterdam, NL, 1994.
- [260] C. Galindo-Legaria, A. Pellenkoft, and M. Kersten. The impact of catalogs and join algorithms on probabilistic query optimization. Technical Report CS-R9459, CWI, Amsterdam, NL, 1994.
- [261] C. Galindo-Legaria, A. Pellenkoft, and M. Kersten. Uniformly-distributed random generation of join orders. Technical Report CS-R9431, CWI, Amsterdam, NL, 1994.
- [262] C. Galindo-Legaria, A. Pellenkoft, and M. Kersten. Cost distribution of search spaces in query optimization. In *Proc. Int. Conf. on Database Theory (ICDT)*, pages 280–293, 1995.
- [263] C. Galindo-Legaria and A. Rosenthal. Outerjoin simplification and re-ordering for query optimization. *ACM Trans. on Database Systems*, 22(1):43–73, Marc 1997.
- [264] C. Galindo-Legaria, A. Rosenthal, and E. Kortright. Expressions, graphs, and algebraic identities for joins, 1991. working paper.

- [265] S. Ganapathy and V. Rajaraman. Information theory applied to the conversion of decision tables to computer programs. *Communications of the ACM*, 16:532–539, 1973.
- [266] S. Gandeharizadeh, J. Stone, and R. Zimmermann. Techniques to quantify SCSI-2 disk subsystem specifications for multimedia. Technical Report 95-610, University of Southern California, 1995.
- [267] S. Ganguly. On the complexity of finding optimal join order sequence for star queries without cross products. personal correspondance, 2000.
- [268] R. Ganski and H. Wong. Optimization of nested SQL queries revisited. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 23–33, 1987.
- [269] G. Garani and R. Johnson. Joining nested relations and subrelations. *Information Systems*, 25(4):287–307, 2000.
- [270] H. Garcia-Molina, J. Ullman, and J. Widom. *Database System Implementation*. Prentice Hall, 2000.
- [271] D. Gardy and L. Nemirovski. Urn models and yao’s formula. In *Proc. Int. Conf. on Database Theory (ICDT)*, pages 100–112, 1999.
- [272] M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, 1978.
- [273] M. R. Garey and D. S. Johnson. *Computers and Intractability: a Guide to the Theory of NP-Completeness*. Freeman, San Francisco, 1979.
- [274] M. Garofalakis and P. Gibbons. Wavelet synopses with error guarantees. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 476–487, 2002.
- [275] P. Gassner, G. Lohman, and K. Schiefer. Query optimization in the IBM DB2 family. *IEEE Data Engineering Bulletin*, 16:4–18, Dec. 1993.
- [276] P. Gassner, G. Lohman, K. Schiefer, and Y. Wang. Query optimization in the IBM DB2 family. Technical report rj 9734, IBM, 1994.
- [277] I. Gent and T. Walsh. Towards an understanding of hill-climbing procedures for SAT. In *Proc. National Conference on Artificial Intelligence*, pages 28–33, 1993.
- [278] P. Gibbons and Y. Matias. New sampling-based summary statistics for improving approximate query answers. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 331–342, 1998.
- [279] P. Godfrey, J. Gryz, and C. Zuzarte. Exploiting constraint-like data characterizations in query optimization. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 582–592, 2001.

- [280] P. Goel and B. Iyer. SLQ query optimization: reordering for a general class of queries. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 47–56, 1996.
- [281] D. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, 1989.
- [282] J. Goldstein and P. Larson. Optimizing queries using materialized views: A practical, scalable solution. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 331–342, 2001.
- [283] J. Goldstein, R. Ramakrishnan, and U. Shaft. Compressing relations and indexes. In *Proc. IEEE Conference on Data Engineering*, 1998. to appear.
- [284] G. Gorry and S. Morton. A framework for management information systems. *Sloan Management Review*, 13(1):55–70, 1971.
- [285] G. Gottlob, C. Koch, and R. Pichler. Efficient algorithms for processing XPath queries. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 95–106, 2002.
- [286] G. Gottlob, C. Koch, and R. Pichler. XPath processing in a nutshell. *SIGMOD Record*, 2003.
- [287] G. Gottlob, C. Koch, and R. Pichler. XPath query evaluation: Improving time and space efficiency. In *Proc. IEEE Conference on Data Engineering*, page to appear, 2003.
- [288] M. Gouda and U. Dayal. Optimal semijoin schedules for query processing in local distributed database systems. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 164–175, 1981.
- [289] P. Goyal. Coding methods for text string search on compressed databases. *Information Systems*, 8(3):231–233, 1983.
- [290] G. Graefe. Software modularization with the exodus optimizer generator. *IEEE Database Engineering*, 9(4):37–45, 1986.
- [291] G. Graefe. Relational division: Four algorithms and their performance. In *Proc. IEEE Conference on Data Engineering*, pages 94–101, 1989.
- [292] G. Graefe. Encapsulation of parallelism in the Volcano query processing system. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages ?–?, 1990.
- [293] G. Graefe. Heap-filter merge join: A new algorithm for joining medium-size inputs. *IEEE Trans. on Software Eng.*, 17(9):979–982, 1991.
- [294] G. Graefe. Query evaluation techniques for large databases. *ACM Computing Surveys*, 25(2), June 1993.
- [295] G. Graefe. Query evaluation techniques for large databases. Shortened version: [294], July 1993.

- [296] G. Graefe. Sort-merge-join: An idea whose time has(h) passed? In *Proc. IEEE Conference on Data Engineering*, pages 406–417, 1994.
- [297] G. Graefe. The cascades framework for query optimization. *IEEE Data Engineering Bulletin*, 18(3):19–29, Sept 1995.
- [298] G. Graefe. Executing nested queries. In *BTW*, pages 58–77, 2003.
- [299] G. Graefe and R. Cole. Fast algorithms for universal quantification in large databases. Internal report, Portland State University and University of Colorado at Boulder, 19??
- [300] G. Graefe and R. Cole. Dynamic query evaluation plans. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages ?–?, 1994.
- [301] G. Graefe and R. Cole. Fast algorithms for universal quantification in large databases. *ACM Trans. on Database Systems*, ?(?) :?–?, ? 1995?
- [302] G. Graefe, R. Cole, D. Davison, W. McKenna, and R. Wolniewicz. Extensible query optimization and parallel execution in Volcano. In *Dagstuhl Query Processing Workshop*, pages 337–380, 1991.
- [303] G. Graefe and D. DeWitt. The EXODUS optimizer generator. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 160–172, 1987.
- [304] G. Graefe, A. Linville, and L. Shapiro. Sort versus hash revisited. *IEEE Trans. on Knowledge and Data Eng.*, 6(6):934–944, Dec. 1994.
- [305] G. Graefe and W. McKenna. The Volcano optimizer generator. Tech. Report 563, University of Colorado, Boulder, 1991.
- [306] G. Graefe and W. McKenna. Extensibility and search efficiency in the volcano optimizer generator. In *Proc. IEEE Conference on Data Engineering*, pages 209–218, 1993.
- [307] G. Graefe and K. Ward. Dynamic query evaluation plans. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 358–366, 1989.
- [308] R. Graham, D. Knuth, and O. Patashnik. *Concrete Mathematics*. Addison-Wesley, 2002.
- [309] G. Grahne and A. Thomo. Algebraic rewritings for optimizing regular path queries. In *Proc. Int. Conf. on Database Theory (ICDT)*, pages 301–315, 2001.
- [310] G. Grahne and A. Thomo. New rewritings and optimizations for regular path queries. In *Proc. Int. Conf. on Database Theory (ICDT)*, pages 242–258, 2003.
- [311] J. Grant, J. Gryz, J. Minker, and L. Raschid. Semantic query optimization for object databases. In *Proc. IEEE Conference on Data Engineering*, pages 444–453, 1997.

- [312] J. Gray, editor. *The Benchmark Handbook*. Morgan Kaufmann Publishers, San Mateo, CA, 1991.
- [313] J. Gray, A. Bosworth, A. Layman, and H. Pirahesh. Data Cube: A relational aggregation operator generalizing group-by, cross-tab, and sub-totals. In *Proc. IEEE Conference on Data Engineering*, pages 152–169, 1996.
- [314] J. Gray and G. Graefe. The five-minute rule ten years later, and other computer storage rules of thumb. *ACM SIGMOD Record*, 26(4):63–68, 1997.
- [315] J. Gray and F. Putzolu. The 5 minute rule for trading memory for disk accesses and the 10 byte rule for trading memory for CPU time. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 395–398, 1987.
- [316] P. Grefen and R. de By. A multi-set extended relational algebra – a formal approach to a practical issue. In *Proc. IEEE Conference on Data Engineering*, pages 80–88, 1994.
- [317] T. Grust. Accelerating XPath location steps. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 109–120, 2002.
- [318] T. Grust and M. Van Keulen. Tree awareness for relational database kernels: Staircase join. In *Intelligent Search on XML Data*, pages 231–245, 2003.
- [319] T. Grust, M. Van Keulen, and J. Teubner. Staircase join: Teach a relational dbms to watch its (axis) steps. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 524–525, 2003.
- [320] T. Grust, M. Van Keulen, and J. Teubner. Accelerating XPath evaluation in any RDBMS. *ACM Trans. on Database Systems*, 29(1):91–131, 2004.
- [321] J. Gryz, B. Schiefer, J. Zheng, and C. Zuzarte. Discovery and application of check constraints in DB2. In *Proc. IEEE Conference on Data Engineering*, pages 551–556, 2001.
- [322] E. Gudes and A. Reiter. On evaluating boolean expression. *Software Practice and Experience*, 3:345–350, 1973.
- [323] H. Gunadhi and A. Segev. Query processing algorithms for temporal intersection joins. In *Proc. IEEE Conference on Data Engineering*, pages 336–344, 1991.
- [324] L. Guo, K. Beyer, J. Shanmugasundaram, and E. Shekita. Efficient inverted lists and query algorithms for structured value ranking in update-intensive relational databases. In *Proc. IEEE Conference on Data Engineering*, pages 298–309, 2005.

- [325] M. Guo, S. Y. W. Su, and H. Lam. An association algebra for processing object-oriented databases. In *Proc. IEEE Conference on Data Engineering*, pages ?–?, 1991.
- [326] A. Gupta, V. Harinarayan, and D. Quass. Aggregate-query processing in data warehousing environments. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 358–369, 1995.
- [327] A. Gupta, V. Harinarayan, and D. Quass. Generalized projections: A powerful approach to aggregation. Technical Report, 1995.
- [328] A. Gupta and I. Mumick. Maintenance of materialized views: problems, techniques and applications. *IEEE Data Engineering Bulletin*, 18(2):3–19, 1995.
- [329] R. Güting, R. Zicari, and D. Choy. An algebra for structured office documents. *ACM Trans. on Information Systems*, 7(4):123–157, 1989.
- [330] R. H. Güting. Geo-relational algebra: A model and query language for geometric database systems. In J. W. Schmidt, S. Ceri, and M. Missikoff, editors, *Proc. of the Intl. Conf. on Extending Database Technology*, pages 506–527, Venice, Italy, Mar 1988. Springer-Verlag, Lecture Notes in Computer Science No. 303.
- [331] R. H. Güting. Second-order signature: A tool for specifying data models, query processing, and optimization. Informatik-Bericht 12/1992, ETH Zürich, 1992.
- [332] L. Haas, W. Chang, G. Lohman, J. McPherson, P. Wilms, G. Lapis, B. Lindsay, H. Pirahesh, M. Carey, and E. Shekita. Starburst mid-flight: As the dust clears. *IEEE Trans. on Knowledge and Data Eng.*, 2(1):143–160, 1990.
- [333] L. Haas, J. Freytag, G. Lohman, and H. Pirahesh. Extensible query processing in starburst. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 377–388, 1989.
- [334] A. Halevy. Answering queries using views: A survey. *The VLDB Journal*, 10(4):270–294, Dec. 2001.
- [335] P. A. V. Hall. Common subexpression identification in general algebraic systems. Tech. rep. uksc 0060, IBM UK Scientific Center, Peterlee, England, 1974.
- [336] P. A. V. Hall. Optimization of single expressions in a relational database system. *IBM J. Res. Devel.*, 20(3):244–257, 1976.
- [337] P. A. V. Hall and S. Todd. Factorization of algebraic expressions. Tech. Report UKSC 0055, IBM UK Scientific Center, Peterlee, England, 1974.
- [338] C. Hamalainen. Complexity of query optimisation and evaluation. Master’s thesis, Griffith University, Queensland, Australia, 2002.

- [339] M. Hammer and B. Niamir. A heuristic approach to attribute partitioning. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 93–101, 1979.
- [340] J. Han. Smallest-first query evaluation for database systems. In *Australian Database Conference*, pages ?–?, Christchurch, New Zealand, Jan. 1994.
- [341] M. Z. Hanani. An optimal evaluation of boolean expressions in an online query system. *Communications of the ACM*, 20(5):344–347, 1977.
- [342] E. Hanson. A performance analysis of view materialization strategies. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 440–453, 1987.
- [343] E. Hanson. Processing queries against database procedures. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages ?–?, 1988.
- [344] E.N. Hanson, M. Chaabouni, C.-H. Kim, and Y.-W. Wang. A predicate matching algorithm for database rule systems. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 271–?, 1990.
- [345] T. Härder. Implementing a generalized access path structure for a relational database system. *ACM Trans. on Database Systems*, 3(3):285–298, 1978.
- [346] T. Härder, B. Mitschang, and H. Schöning. Query processing for complex objects. *Data and Knowledge Engineering*, 7(3):181–200, 1992.
- [347] T. Härder and E. Rahm. *Datenbanksysteme*. Springer, 1999.
- [348] T. Härder, H. Schöning, and A. Sikeler. Parallelism in processing queries on complex objects. In *International Symposium on Databases in Parallel and Distributed Systems*, Ausgin, TX, August 1988.
- [349] V. Harinarayan. Issues in interactive aggregation. *IEEE Data Engineering Bulletin*, 20(1):12–18, 1997.
- [350] V. Harinarayan and A. Gupta. Generalized projections: a powerful query optimization technique. Technical Report STAN-CS-TN-94-14, Stanford University, 1994.
- [351] E. Harris and K. Ramamohanarao. Join algorithm costs revisited. *The VLDB Journal*, 5(1):?–?, Jan 1996.
- [352] W. Hasan and H. Pirahesh. Query rewrite optimization in starburst. Research Report RJ6367, IBM, 1988.
- [353] Heller. Rabbit: A performance counter library for Intel/AMD processors and Linux. perform internet search for this or similar tools.
- [354] J. Hellerstein. Practical predicate placement. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 325–335, 1994.

- [355] J. Hellerstein and J. Naughton. Query execution techniques for caching expensive methods. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 423–434, 1996.
- [356] J. Hellerstein and M. Stonebraker. Predicate migration: Optimizing queries with expensive predicates. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 267–277, 1993.
- [357] S. Helmer, C.-C. Kanne, and G. Moerkotte. Optimized translation of XPath expressions into algebraic expressions parameterized by programs containing navigational primitives. In *Proc. Int. Conf. on Web Information Systems Engineering (WISE)*, 2002. 215-224.
- [358] S. Helmer, B. König-Ries, and G. Moerkotte. The relational difference calculus and applications. Technical report, Universität Karlsruhe, 1993. (unpublished manuscript).
- [359] S. Helmer and G. Moerkotte. Evaluation of main memory join algorithms for joins with set comparison join predicates. Technical Report 13/96, University of Mannheim, Mannheim, Germany, 1996.
- [360] S. Helmer and G. Moerkotte. Evaluation of main memory join algorithms for joins with set comparison join predicates. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 386–395, 1997.
- [361] S. Helmer and G. Moerkotte. Index structures for databases containing data items with set-valued attributes. Technical Report 2/97, University of Mannheim, 1997.
- [362] S. Helmer and G. Moerkotte. A study of four index structures for set-valued attributes of low cardinality. Technical Report 02/99, University of Mannheim, 1999.
- [363] S. Helmer and G. Moerkotte. Compiling away set containment and intersection joins. Technical Report 4, University of Mannheim, 2002.
- [364] S. Helmer and G. Moerkotte. A performance study of four index structures for set-valued attributes of low cardinality. *VLDB Journal*, 12(3):244–261, 2003.
- [365] S. Helmer, T. Neumann, and G. Moerkotte. Early grouping gets the skew. Technical Report 9, University of Mannheim, 2002.
- [366] S. Helmer, T. Neumann, and G. Moerkotte. Estimating the output cardinality of partial preaggregation with a measure of clusteredness. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 656–667, 2003.
- [367] A. Heuer and M. H. Scholl. Principles of object-oriented query languages. In *Proc. der GI-Fachtagung Datenbanksysteme für Büro, Technik und Wissenschaft (BTW)*. Springer, 1991.

- [368] J. Hidders and P. Michiels. Avoiding unnecessary ordering operations in XPath. In *Int. Workshop on Database Programming Languages*, pages 54–70, 2003.
- [369] D. Hirschberg. On the complexity of searching a set of vectors. *SIAM J. Computing*, 9(1):126–129, 1980.
- [370] T. Hogg and C. Williams. Solving the really hard problems with cooperative search. In *Proc. National Conference on Artificial Intelligence*, pages 231–236, 1993.
- [371] L. Hong-Cheu and K. Ramamohanarao. Algebraic equivalences among nested relational expressions. In *CIKM*, pages 234–243, 1994.
- [372] V. Hristidis, N. Koudas, and Y. Papakonstantinou. PREFER: A system for the efficient execution of multi-parametric ranked queries. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages ?–?, 2001.
- [373] N. Huyn. Multiple-view self-maintenance in data warehousing environments. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 26–35, 1997.
- [374] F. Hwang and G. Chang. Enumerating consecutive and nested partitions for graphs. Technical Report DIMACS Technical Report 93-15, Rutgers University, 1993.
- [375] H.-Y. Hwang and Y.-T. Yu. An analytical method for estimating and interpreting query time. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 347–358, 1987.
- [376] L. Hyafil and R. Rivest. Constructing optimal binary decision trees is NP-complete. *Information Processing Letters*, 5(1):15–17, 1976.
- [377] T. Ibaraki and T. Kameda. Optimal nesting for computing n-relational joins. *ACM Trans. on Database Systems*, 9(3):482–502, 1984.
- [378] O. Ibarra and J. Su. On the containment and equivalence of database queries with linear constraints. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages 32–43, 1997.
- [379] A. IJbema and H. Blanken. Estimating bucket accesses: A practical approach. In *Proc. IEEE Conference on Data Engineering*, pages 30–37, 1986.
- [380] I. Ilyas, J. Rao, G. Lohman, D. Gao, and E. Lin. Estimating compilation time of a query optimizer. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 373–384, 2003.
- [381] Y. Ioannidis. Query optimization. *ACM Computing Surveys*, 28(1):121–123, 1996.

- [382] Y. Ioannidis. *A. Tucker (ed.): The Computer Science and Engineering Handbook*, chapter Query Optimization, pages 1038–1057. CRC Press, 1997.
- [383] Y. Ioannidis and R. Ramakrishnan. Containment of conjunctive queries: Beyond relations and sets. *ACM Trans. on Database Systems*, 20(3):288–324, 1992.
- [384] Y. E. Ioannidis and S. Christodoulakis. On the propagation of errors in the size of join results. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 268–277, 1991.
- [385] Y. E. Ioannidis and Y. C. Kang. Randomized algorithms for optimizing large join queries. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 312–321, 1990.
- [386] Y. E. Ioannidis and Y. C. Kang. Left-deep vs. bushy trees: An analysis of strategy spaces and its implications for query optimization. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 168–177, 1991.
- [387] Y. E. Ioannidis, Y. C. Kang, and T. Zhang. Cost wells in random graphs. personal communication, Dec. 1996.
- [388] Y. E. Ioannidis, R. T. Ng, K. Shim, and T. K. Sellis. Parametric query optimization. Tech. report, University of Wisconsin, Madison, 1992.
- [389] Y. E. Ioannidis, R. T. Ng, K. Shim, and T. K. Sellis. Parametric query optimization. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 103–114, 1992.
- [390] Y. E. Ioannidis and E. Wong. Query optimization by simulated annealing. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 9–22, 1987.
- [391] D. Jacobsen and J. Wilkes. Disk scheduling algorithms based on rotational position. Technical Report HPL-CSP-91-7, Hewlett-Packard Laboratories, 1991.
- [392] H. V. Jagadish, S. Al-Khalifa, A. Chapman, L.V.S. Lakshmanan, A. Nierman, S. Papparizos, J. Patel, D. Srivastava, N. Wiwatwattana, Y. Wu, and C. Yu. TIMBER: A Native XML Database. *VLDB Journal*, 2003. to appear.
- [393] H. V. Jagadish, L. V. S. Lakshmanan, D. Srivastava, and K. Thompson. TAX: A tree algebra for XML. In *Proc. Int. Workshop on Database Programming Languages*, pages 149–164, 2001.
- [394] C. Janssen. The visual profiler. perform internet search for this or similar tools.

- [395] M. Jarke. Common subexpression isolation in multiple query optimization. In W. Kim, D. Reiner, and D. Batory, editors, *Topics in Information Systems. Query Processing in Database Systems*, pages 191–205, 1985.
- [396] M. Jarke. Common subexpression isolation in multiple query optimization. In *Query Processing in Database Systems, W. Kim, D. Reiner, D. Batory (Eds.)*, pages 191–205, 1985.
- [397] M. Jarke and J.Koch. Range nesting: A fast method to evaluate quantified queries. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 196–206, 1983.
- [398] M. Jarke and J. Koch. Query optimization in database systems. *ACM Computing Surveys*, pages 111–152, Jun 1984.
- [399] P. Jenq, D. Woelk, W. Kim, and W. Lee. Query processing in distributed ORION. In *Proc. of the Int. Conf. on Extending Database Technology (EDBT)*, pages 169–187, 1990.
- [400] A. Jhingran. A performance study of query optimization algorithms on a database system supporting procedures. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 88–99, 1988.
- [401] D. Johnson and A. Klug. Testing containment of conjunctive queries under functional and inclusion dependencies. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages 164–168, 1982.
- [402] D. Johnson and A. Klug. Testing containment of conjunctive queries under functional and inclusion dependencies. *J. Comp. Sys. Sci.*, 28(1):167–189, 1984.
- [403] D. S. Johnson and A. Klug. Optimizing conjunctive queries that contain untyped variables. *SIAM J. Comput.*, 12(4):616–640, 1983.
- [404] B. Jonsson, M. Franklin, and D. Srivastava. Interaction of query evaluation and buffer management for information retrieval. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 118–129, 1998.
- [405] N. Kabra and D. DeWitt. Efficient mid-query re-optimization of sub-optimal query execution plans. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 106–117, 1998.
- [406] J. Kahn, G. Kalai, and N. Linial. The influence of variables on boolean functions. In *IEEE ???*, pages 68–80, 1988.
- [407] M. Kamath and K. Ramamritham. Bucket skip merge join: A scalable algorithm for join processing in very large databases using indexes. Technical Report 20, University of Massachusetts at Amherst, Amherst, MA, 1996.

- [408] Y. Kambayashi. Processing cyclic queries. In W. Kim, D. Reiner, and D. Batory, editors, *Query Processing in Database Systems*, pages 62–78, 1985.
- [409] Y. Kambayashi and M. Yoshikawa. Query processing utilizing dependencies and horizontal decomposition. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 55–68, 1983.
- [410] A. Kawaguchi, D. Lieuwen, I. Mumick, and K. Ross. Implementing incremental view maintenance in nested data models. In *Proc. Int. Workshop on Database Programming Languages*, 1997.
- [411] A. Keller and J. Basu. A predicate-based caching scheme for client-server database architectures. In *PDIS*, pages 229–238, 1994.
- [412] T. Keller, G. Graefe, and D. Maier. Efficient assembly of complex objects. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 148–157, 1991.
- [413] A. Kemper and A. Eickler. *Datenbanksysteme*. Oldenbourg, 2001. 4th Edition.
- [414] A. Kemper and G. Moerkotte. Advanced query processing in object bases: A comprehensive approach to access support, query transformation and evaluation. Technical Report 27/90, University of Karlsruhe, 1990.
- [415] A. Kemper and G. Moerkotte. Advanced query processing in object bases using access support relations. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 294–305, 1990.
- [416] A. Kemper and G. Moerkotte. Query optimization in object bases: Exploiting relational techniques. In *Proc. Dagstuhl Workshop on Query Optimization (J.-C. Freytag, D. Maier und G. Vossen (eds.))*. Morgan-Kaufman, 1993.
- [417] A. Kemper, G. Moerkotte, and K. Peithner. A blackboard architecture for query optimization in object bases. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 543–554, 1993.
- [418] A. Kemper, G. Moerkotte, K. Peithner, and M. Steinbrunn. Optimizing disjunctive queries with expensive predicates. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 336–347, 1994.
- [419] A. Kemper, G. Moerkotte, and M. Steinbrunn. Optimierung Boolescher Ausdrücke in Objektbanken. In *Grundlagen von Datenbanken (Eds. U. Lipeck, R. Manthey)*, pages 91–95, 1992.
- [420] A. Kemper, G. Moerkotte, and M. Steinbrunn. Optimization of boolean expressions in object bases. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 79–90, 1992.

- [421] W. Kiessling. SQL-like and Quel-like correlation queries with aggregates revisited. ERL/UCB Memo 84/75, University of Berkeley, 1984.
- [422] W. Kiessling. On semantic reefs and efficient processing of correlation queries with aggregates. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 241–250, 1985.
- [423] K. C. Kim, W. Kim, D. Woelk, and A. Dale. Acyclic query processing in object-oriented databases. In *Proc. of the Entity Relationship Conf.*, 1988.
- [424] W. Kim. On optimizing an SQL-like nested query. *ACM Trans. on Database Systems*, 7(3):443–469, Sep 82.
- [425] J. J. King. Exploring the use of domain knowledge for query processing efficiency. Technical Report STAN-CS-79-781, Computer Science Department, Stanford University, 1979.
- [426] J. J. King. Quist: A system for semantic query optimization in relational databases. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 510–517, 1981.
- [427] A. Klausner. *Multirelations in Relational Databases*. PhD thesis, Harvard University, Cambridge, 1986.
- [428] M. Klettke, L. Schneider, and A. Heuer. Metrics for XML Document Collections. In *EDBT Workshop XML-Based Data Management (XMLDM)*, pages 15–28, 2002.
- [429] A. Klug. Calculating constraints on relational expressions. *ACM Trans. on Database Systems*, 5(3):260–290, 1980.
- [430] A. Klug. Access paths in the “ABE” statistical query facility. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 161–173, 1982.
- [431] A. Klug. Equivalence of relational algebra and relational calculus query languages having aggregate functions. *Journal of the ACM*, 29(3):699–717, 1982.
- [432] A. Klug. On conjunctive queries containing inequalities. *Journal of the ACM*, 35(1):146–160, 1988. Written 1982 and published posthumously.
- [433] D. Knuth. *The Art of Computer Programming, Volume 1: Fundamental Algorithms*. Addison-Wesley, 1997.
- [434] J. Koch. *Relationale Anfragen: Zerlegung und Optimierung*. Informatik-Fachberichte 101. Springer-Verlag, 1985.
- [435] J. Kollias. An estimate for seek time for batched searching of random or index sequential structured files. *The Computer Journal*, 21(2):132–133, 1978.

- [436] A. König and G. Weikum. Combining histograms and parametric curve fitting for feedback-driven query result-size estimation. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 423–434, 1999.
- [437] B. König-Ries, S. Helmer, and G. Moerkotte. An experimental study on the complexity of left-deep join ordering problems for cyclic queries. Working Draft, 1994.
- [438] B. König-Ries, S. Helmer, and G. Moerkotte. An experimental study on the complexity of left-deep join ordering problems for cyclic queries. Technical Report 95-4, RWTH-Aachen, 1995.
- [439] D. Kossmann. The state of the art in distributed query processing. *ACM Computing Surveys*, 32(4):422–469, 2000.
- [440] D. Kossmann and K. Stocker. Iterative dynamic programming: a new class of query optimization algorithms. *ACM Trans. on Database Systems*, 25(1):43–82, 2000.
- [441] W. Kowarschick. Semantic optimization: What are disjunctive residues useful for? *SIGMOD Record*, 21(3):26–32, September 1992.
- [442] D. Kreher and D. Stinson. *Combinatorial Algorithms: Generation, Enumeration, and Search*. CRC Press, 1999.
- [443] R. Krishnamurthy, H. Boral, and C. Zaniolo. Optimization of nonrecursive queries. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 128–137, 1986.
- [444] A. Kumar and M. Stonebraker. The effect of join selectivities on optimal nesting order. *SIGMOD Record*, 16(1):28–41, 1987.
- [445] I. Kunen and D. Suciu. A scalable algorithm for query minimization. ask Dan for more information, year.
- [446] S. Kwan and H. Strong. Index path length evaluation for the research storage system of system r. Technical Report RJ2736, IBM Research Laboratory, San Jose, 1980.
- [447] L. Lakshman and R. Missaoui. Pushing semantics inside recursion: A general framework for semantic optimization of recursive queries. In *Proc. IEEE Conference on Data Engineering*, pages 211–220, 1995.
- [448] S. Lang and Y. Manolopoulos. Efficient expressions for completely and partly unsuccessful batched search of tree-structured files. *IEEE Trans. on Software Eng.*, 16(12):1433–1435, 1990.
- [449] S.-D. Lang, J. Driscoll, and J. Jou. A unified analysis of batched searching of sequential and tree-structured files. *ACM Trans. on Database Systems*, 14(4):604–618, 1989.

- [450] T. Lang, C. Wood, and I. Fernandez. Database buffer paging in virtual storage systems. *ACM Trans. on Database Systems*, 2(4):339–351, 1977.
- [451] R. Lanzelotte and J.-P. Cheiney. Adapting relational optimisation technology for deductive and object-oriented declarative database languages. In *Proc. Int. Workshop on Database Programming Languages*, pages 322–336, 1991.
- [452] R. Lanzelotte and P. Valduriez. Extending the search strategy in a query optimizer. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 363–373, 1991.
- [453] R. Lanzelotte, P. Valduriez, and M. Zait. Optimization of object-oriented recursive queries using cost-controlled strategies. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 256–265, 1992.
- [454] R. Lanzelotte, P. Valduriez, and M. Zait. On the effectiveness of optimization search strategies for parallel execution. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 493–504, 1993.
- [455] R. Lanzelotte, P. Valduriez, M. Ziane, and J.-P. Cheiney. Optimization of nonrecursive queries in OODBMs. In *Proc. Int. Conf. on Deductive and Object-Oriented Databases (DOOD)*, pages 1–21, 1991.
- [456] P.-A. Larson. Data reduction by partial preaggregation. In *Proc. IEEE Conference on Data Engineering*, pages 706–715, 2002.
- [457] P.-Å. Larson and H. Yang. Computing queries from derived relations. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 259–269, 1985.
- [458] Y.-N. Law, H. Wang, and C. Zaniolo. Query languages and data models for database sequences and data streams. In *VLDB*, pages 492–503, 2004.
- [459] E. Lawler. Sequencing jobs to minimize total weighted completion time subject to precedence constraints. *Ann. Discrete Math.*, 2:75–90, 1978.
- [460] C. Lee, C.-S. Shih, and Y.-H. Chen. Optimizing large join queries using a graph-based approach. *IEEE Trans. on Knowledge and Data Eng.*, 13(2):298–315, 2001.
- [461] M. K. Lee, J. C. Freytag, and G. M. Lohman. Implementing an interpreter for functional rules in a query optimizer. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 218–239, 1988.
- [462] M. K. Lee, J. C. Freytag, and G. M. Lohman. Implementing an optimizer for functional rules in a query optimizer. Technical Report RJ 6125, IBM Almaden Research Center, San Jose, CA, 1988.
- [463] M.K. Lee, J.C. Freytag, and G.M. Lohman. Implementing an interpreter for functional rules in a query optimizer. Research report RJ 6125, IBM, 1988.

- [464] T. Lehman and B. Lindsay. The Starburst long field manager. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 375–383, 1989.
- [465] K. Lehnert. *Regelbasierte Beschreibung von Optimierungsverfahren für relationale Datenbankanfragesprachen*. PhD thesis, Technische Universität München, 8000 München, West Germany, Dec 1988.
- [466] A. Lerner and D. Shasha. AQuery: query language for ordered data, optimization techniques, and experiments. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 345–356, 2003.
- [467] H. Leslie, R. Jain, D. Birdsall, and H. Yaghmai. Efficient search of multi-dimensional B-trees. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 710–719, 1995.
- [468] C. Leung, H. Pirahesh, and P. Seshadri. Query rewrite optimization rules in IBM DB2 universal database. Research Report RJ 10103 (91919), IBM Almaden Research Division, January 1998.
- [469] M. Levene and G. Loizou. Correction to null values in nested relational databases by m. roth and h. korth and a. silberschatz. *Acta Informatica*, 28(6):603–605, 1991.
- [470] M. Levene and G. Loizou. A fully precise null extended nested relational algebra. *Fundamenta Informaticae*, 19(3/4):303–342, 1993.
- [471] A. Levy, A. Mendelzon, and Y. Sagiv. Answering queries using views. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages ?–?, 1995.
- [472] A. Levy, A. Mendelzon, Y. Sagiv, and D. Srivastava. *Answering Queries Using Views*, pages 93–106. MIT Press, 1999.
- [473] A. Levy, A. Mendelzon, D. Srivastava, and Y. Sagiv. Answering queries using views. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages 95–104, 1995.
- [474] A. Levy, I. Mumick, and Y. Sagiv. Query optimization by predicate move-around. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 96–107, 1994.
- [475] A.Y. Levy and I.S. Mumick. Reasoning with aggregation constraints. In P. Apers, M. Bouzeghoub, and G. Gardarin, editors, *Proc. European Conf. on Extending Database Technology (EDBT)*, Lecture Notes in Computer Science, pages 514–534. Springer, March 1996.
- [476] H. Lewis and C. Papadimitriou. *Elements of the Theory of Computation*. Prentice Hall, 1981.
- [477] C. Li, K. Chang, I. Ilyas, and S. Song. RankSQL: Query algebra and optimization for relational top-k queries. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 131–142, 2005.

- [478] D. Lichtenstein. Planar formulae and their uses. *SIAM J. Comp.*, 11(2):329–343, 1982.
- [479] J. Liebehenschel. Ranking and unranking of lexicographically ordered words: An average-case analysis. *J. of Automata, Languages, and Combinatorics*, 2:227–268, 1997.
- [480] J. Liebehenschel. Lexicographical generation of a generalized dyck language. Technical Report 5/98, University of Frankfurt, 1998.
- [481] J. Liebehenschel. *Lexikographische Generierung, Ranking und Unranking kombinatorischer Objekt: Eine Average-Case Analyse*. PhD thesis, University of Frankfurt, 2000.
- [482] H. Liefke. Horizontal query optimization on ordered semistructured data. In *ACM SIGMOD Workshop on the Web and Databases (WebDB)*, 1999.
- [483] L. Lim, M. Wang, S. Padmanabhan, J. Vitter, and R. Parr. XPathLearner: An on-line self-tuning Markov histogram for XML path selectivity estimation. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 442–453, 2002.
- [484] J. Lin and M. Ozsoyoglu. Processing OODB queries by O-algebra. In *Int. Conference on Information and Knowledge Management (CIKM)*, pages 134–142, 1996.
- [485] J. W. S. Liu. Algorithms for parsing search queries in systems with inverted file organization. *ACM Trans. on Database Systems*, 1(4):299–316, 1976.
- [486] M.-L. Lo and C. Ravishankar. Towards eliminating random I/O in hash joins. In *Proc. IEEE Conference on Data Engineering*, pages 422–429, 1996.
- [487] G. Lohman. Grammar-like functional rules for representing query optimization alternatives. Research report rj 5992, IBM, 1987.
- [488] G. M. Lohman. Grammar-like functional rules for representing query optimization alternatives. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 18–27, 1988.
- [489] D. Lomet. B-tree page size when caching is considered. *ACM SIGMOD Record*, 27(3):28–32, 1998.
- [490] R. Lorie. XRM - an extended (N-ary) relational model. Technical Report 320-2096, IBM Cambridge Scientific Center, 1974.
- [491] H. Lu and K.-L. Tan. On sort-merge algorithms for band joins. *IEEE Trans. on Knowledge and Data Eng.*, 7(3):508–510, Jun 1995.
- [492] W. S. Luk. On estimating block accesses in database organizations. *Communications of the ACM*, 26(11):945–947, 1983.

- [493] G. Luo, J. Naughton, C. Ellmann, and M. Watzke. Toward a progress indicator for database queries. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 791–802, 2004.
- [494] G. Luo, J. Naughton, C. Ellmann, and M. Watzke. Increasing the accuracy and coverage of SQL progress indicators. In *Proc. IEEE Conference on Data Engineering*, pages 853–864, 2005.
- [495] D. Maier and D. S. Warren. Incorporating computed relations in relational databases. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 176–187, 1981.
- [496] A. Makinouchi, M. Tezuka, H. Kitakami, and S. Adachi. The optimization strategy for query evaluation in RDB/V1. In *Proc. IEEE Conference on Data Engineering*, pages 518–529, 1981.
- [497] T. Malkemus, S. Padmanabhan, and B. Bhattacharjee. Predicate derivation and monotonicity detection in DB2 UDB. In *Proc. IEEE Conference on Data Engineering*, pages ?–?, 2005.
- [498] C. V. Malley and S. B. Zdonik. A knowledge-based approach to query optimization. In *Proc. Int. Conf. on Expert Database Systems*, pages 329–344, 1987.
- [499] N. Mamoulis. Efficient processing of joins on set-valued attributes. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 157–168, 2003.
- [500] G. Manku, S. Rajagopalan, and B. Lindsay. Approximate medians and other quantiles in one pass and with limited memory. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 426–435, Seattle, WA, 1998.
- [501] M. V. Mannino, P. Chu, and T. Sager. Statistical profile estimation in database systems. *ACM Computing Surveys*, 20(3):191–221, 1988.
- [502] Y. Manolopoulos and J. Kollias. Estimating disk head movement in batched searching. *BIT*, 28:27–36, 1988.
- [503] Y. Manolopoulos, J. Kollias, and M. Hatzopoulos. Sequential vs. binary batched search. *The Computer Journal*, 29(4):368–372, 1986.
- [504] Y. Manolopoulos and J. Kollias. Performance of a two-headed disk system when serving database queries under the scan policy. *ACM Trans. on Database Systems*, 14(3):425–442, 1989.
- [505] S. March and D. Severence. The determination of efficient record segmentation and blocking factors for shared data files. *ACM Trans. on Database Systems*, 2(3):279–296, 1977.
- [506] R. Marek and E. Rahm. TID hash joins. In *Int. Conference on Information and Knowledge Management (CIKM)*, pages 42–49, 1994.

- [507] N. May, S. Helmer, C.-C. Kanne, and G. Moerkotte. Xquery processing in natix with an emphasis on join ordering. In *Int. Workshop on XQuery Implementation, Experience and Perspectives (XIME-P)*, pages 49–54, 2004.
- [508] N. May, S. Helmer, and G. Moerkotte. Nested queries and quantifiers in an ordered context. Technical report, University of Mannheim, 2003.
- [509] N. May, S. Helmer, and G. Moerkotte. Quantifiers in XQuery. In *Proc. Int. Conf. on Web Information Systems Engineering (WISE)*, pages 313–316, 2003.
- [510] N. May, S. Helmer, and G. Moerkotte. Three Cases for Query Decorrelation in XQuery. In *Int. XML Database Symp. (XSym)*, pages 70–84, 2003.
- [511] N. May, S. Helmer, and G. Moerkotte. Nested queries and quantifiers in an ordered context. In *Proc. IEEE Conference on Data Engineering*, pages 239–250, 2004.
- [512] J. McHugh and J. Widom. Query optimization for XML. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 315–326, 1999.
- [513] N. Megiddo and D. Modha. Outperforming LRU with an adaptive replacement cache algorithm. *IEEE Computer*, 37(4):58–65, 2004.
- [514] S. Melnik and H. Garcia-Molina. Divide-and-conquer algorithm for computing set containment joins. In *Proc. of the Int. Conf. on Extending Database Technology (EDBT)*, pages 427–444, 2002.
- [515] S. Melnik and H. Garcia-Molina. Adaptive algorithms for set containment joins. *ACM Trans. on Database Systems*, 28(1):56–99, 2003.
- [516] T. H. Merrett, Y. Kambayashi, and H. Yasuura. Scheduling of page-fetches in join operations. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 488–498, 1981.
- [517] R. Van Meter. Observing the effects of multi-zone disks. In *USENIX Annual Technical Conference*, 1997.
- [518] G. Miklau and D. Suciu. Containment and equivalence of XPath expressions. *Journal of the ACM*, 51(1):2–45, 2002.
- [519] T. Milo and D. Suciu. Index structures for path expressions. In *Proc. Int. Conf. on Database Theory (ICDT)*, pages 277–295, 1999.
- [520] M. Minoux. *Mathematical Programming. Theory and Algorithms*. Wiley, 1986.
- [521] D. Mitchell, B. Selman, and H. Levesque. Hard and easy distributions of SAT problems. In *Proc. National Conference on Artificial Intelligence*, pages 459–465, 1992.

- [522] G. Mitchell. *Extensible Query Processing in an Object-Oriented Database*. PhD thesis, Brown University, Providence, RI 02912, 1993.
- [523] G. Mitchell, U. Dayal, and S. Zdonik. Control of an extensible query optimizer: A planning-based approach. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages ?–?, 1993.
- [524] G. Mitchell, S. Zdonik, and U. Dayal. Object-oriented query optimization: What’s the problem? Technical Report CS-91-41, Brown University, 1991.
- [525] G. Mitchell, S. Zdonik, and U. Dayal. An architecture for query processing in persistent object stores. In *Proc. of the Hawaiian Conf. on Computer and System Sciences*, pages 787–798, 1992.
- [526] G. Mitchell, S. Zdonik, and U. Dayal. *A. Dogac and M. T. Özsu and A. Biliris, and T. Sellis: Object-Oriented Database Systems*, chapter Optimization of Object-Oriented Queries: Problems and Applications, pages 119–146. NATO ASI Series F: Computer and Systems Sciences, Vol. 130. Springer, 1994.
- [527] G. Moerkotte. Small materialized aggregates: A light weight index structure for data warehousing. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 476–487, 1998.
- [528] G. Moerkotte. Constructing Optimal Bushy Trees Possibly Containing Cross Products for Order Preserving Joins is in P. Technical Report 12, University of Mannheim, 2003.
- [529] G. Moerkotte. Dp-counter analytics. Technical Report 2, University of Mannheim, 2006.
- [530] G. Moerkotte and T. Neumann. Analysis of two existing and one new dynamic programming algorithm for the generation of optimal bushy trees without cross products. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 930–941, 2006.
- [531] C. Mohan. Interactions between query optimization and concurrency control. In *Int. Workshop on RIDE*, 1992.
- [532] C. Mohan, D. Haderle, Y. Wang, and J. Cheng. Single table access using multiple indexes: Optimization, execution, and concurrency control techniques. In *Int. Conf. on Extended Database Technology (EDBT)*, pages 29–43, 1990.
- [533] C. Monma and J. Sidney. Sequencing with series-parallel precedence constraints. *Math. Oper. Res.*, 4:215–224, 1979.
- [534] T. Morzy, M. Matyasiak, and S. Salza. Tabu search optimization of large join queries. In *Proc. of the Int. Conf. on Extending Database Technology (EDBT)*, pages 309–322, 1994.

- [535] L. Moses and R. Oakland. *Tables of Random Permutations*. Stanford University Press, 1963.
- [536] I. Mumick, S. Finkelstein, H. Pirahesh, and R. Ramakrishnan. Magic is relevant. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 247–258, 1990.
- [537] I. Mumick and H. Pirahesh. Implementation of magic sets in a relational database system. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 103–114, 1994.
- [538] I. Mumick, H. Pirahesh, and R. Ramakrishnan. The magic of duplicates and aggregates. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 264–277, 1990.
- [539] M. Muralikrishna. Optimization and dataflow algorithms for nested tree queries. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, 1989.
- [540] M. Muralikrishna. Improved unnesting algorithms for join aggregate SQL queries. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 91–102, 1992.
- [541] W. Myrvold and F. Ruskey. Ranking and unranking permutations in linear time. *Information Processing Letters*, 79(6):281–284, 2001.
- [542] R. Nakano. Translation with optimization from relational calculus to relational algebra having aggregate functions. *ACM Trans. on Database Systems*, 15(4):518–557, 1990.
- [543] T. Neumann and G. Moerkotte. A combined framework for grouping and order optimization. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 960–971, 2004.
- [544] T. Neumann and G. Moerkotte. An efficient framework for order optimization. In *Proc. IEEE Conference on Data Engineering*, pages 461–472, 2004.
- [545] F. Neven and T. Schwentick. XPath containment in the presence of disjunction, dtlds, and variables. In *Proc. Int. Conf. on Database Theory (ICDT)*, pages 315–329, 2003.
- [546] S. Ng. Advances in disk technology: Performance issues. *IEEE Computer*, 31(5):75–81, 1998.
- [547] W. Ng and C. Ravishankar. Relational database compression using augmented vector quantization. In *Proc. IEEE Conference on Data Engineering*, pages 540–549, 1995.
- [548] S. Nigam and K. Davis. A semantic query optimization algorithm for object-oriented databases. In *Second International Workshop on Constraint Database Systems*, pages 329–344, 1997.

- [549] E. Omicinski. Heuristics for join processing using nonclustered indexes. *IEEE Trans. on Software Eng.*, 15(1):18–25, Feb. 1989.
- [550] P. O’Neil. *Database Principles, Programming, Performance*. Morgan Kaufmann, 1994.
- [551] P. O’Neil and D. Quass. Improved query performance with variant indexes. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 38–49, 1997.
- [552] K. Ono and G. Lohman. Extensible enumeration of feasible joins for relational query optimization. Technical Report RJ 6625, IBM Almaden Research Center, 1988.
- [553] K. Ono and G. Lohman. Measuring the complexity of join enumeration in query optimization. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 314–325, 1990.
- [554] J. A. Orenstein, S. Haradhvala, B. Margulies, and D. Sakahara. Query processing in the ObjectStore database system. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 403–412, 1992.
- [555] J. A. Orenstein and F. A. Manola. PROBE spatial data modeling and query processing in an image database application. *IEEE Trans. on Software Eng.*, 14(5):611–629, 1988.
- [556] M. Ortega-Binderberger, K. Chakrabarti, and S. Mehrotra. An approach to integrating query refinement in sql. In *Proc. of the Int. Conf. on Extending Database Technology (EDBT)*, pages 15–33, 2002.
- [557] S. Osborn. Identity, equality and query optimization. In *Proc. OODB*, 1989.
- [558] N. Ott and K. Horlaender. Removing redundant joins in queries involving views. Technical Report TR-82.03.003, IBM Heidelberg Scientific Center, Heidelberg, 1982.
- [559] G. Ozsoyoglu, V. Matos, and Z. M. Ozsoyoglu. Query processing techniques in the Summary-Table-by-Example database query language. *ACM Trans. on Database Systems*, 14(4):526–573, 1989.
- [560] G. Özsoyoglu, Z. M. Özsoyoglu, and V. Matos. Extending relational algebra and relational calculus with set-valued attributes and aggregate functions. *ACM Trans. on Database Systems*, 12(4):566–592, Dec 1987.
- [561] G. Ozsoyoglu and H. Wang. A relational calculus with set operators, its safety and equivalent graphical languages. *IEEE Trans. on Software Eng.*, SE-15(9):1038–1052, 1989.
- [562] T. Özsu and J. Blakeley. *W. Kim (ed.): Modern Database Systems*, chapter Query Processing in Object-Oriented Database Systems, pages 146–174. Addison Wesley, 1995.

- [563] T. Özsu and D. Meechan. Finding heuristics for processing selection queries in relational database systems. *Information Systems*, 15(3):359–373, 1990.
- [564] T. Özsu and D. Meechan. Join processing heuristics in relational database systems. *Information Systems*, 15(4):429–444, 1990.
- [565] T. Özsu and P. Valduriez. *Principles of Distributed Database Systems*. Prentice-Hall, 1999.
- [566] T. Özsu and B. Yao. Evaluation of DBMSs using Xbench benchmark. Technical Report CS-2003-24, University of Waterloo, 2003.
- [567] P. Palvia. Expressions for batched searching of sequential and hierarchical files. *ACM Trans. on Database Systems*, 10(1):97–106, 1985.
- [568] P. Palvia and S. March. Approximating block accesses in database organizations. *Information Processing Letters*, 19:75–79, 1984.
- [569] S. Papadimitriou, H. Kitagawa, P. Gibbons, and C. Faloutsos. LOCI: Fast outlier detection using local correlation integral. In *ICDE*, pages 315–, 2003.
- [570] V. Papadimos and D. Maier. Mutant query plans. *Information & Software Technology*, 44(4):197–206, 2002.
- [571] Y. Papakonstantinou and V. Vianu. Incremental validation of XML documents. In *Proc. Int. Conf. on Database Theory (ICDT)*, pages 47–63, 2003.
- [572] S. Paparizos, S. Al-Khalifa, H. V. Jagadish, L. V. S. Lakshmanan, A. Nierman, D. Srivastava, and Y. Wu. Grouping in XML. In *EDBT Workshops*, pages 128–147, 2002.
- [573] S. Paparizos, S. Al-Khalifa, H. V. Jagadish, A. Niermann, and Y. Wu. A physical algebra for XML. Technical report, University of Michigan, 2002.
- [574] J. Paredaens and D. Van Gucht. Converting nested algebra expressions into flat algebra expressions. *ACM Trans. on Database Systems*, 17(1):65–93, Mar 1992.
- [575] C.-S. Park, M. Kim, and Y.-J. Lee. Rewriting OLAP queries using materialized views and dimension hierarchies in data. In *Proc. IEEE Conference on Data Engineering*, pages 515–523, 2001.
- [576] J. Patel, M. Carey, and M. Vernon. Accurate modeling of the hybrid hash join algorithm. In *Proc. ACM SIGMETRICS Conf. on Measurement and Modeling of Computer Systems*, pages 56–66, 1994.
- [577] R. Patterson, G. Gibson, E. Ginting, D. Stodolsky, and J. Zelenka. Informed prefetching and caching. Technical Report CMU-CS-95-134, Carnegie Mellon University, 1995.

- [578] R. Patterson, G. Gibson, and M. Sayanarayanan. A status report on research in transparent informed prefetching. Technical Report CMU-CS-93-113, Carnegie Mellon University, 1993.
- [579] G. Paulley and P.-A. Larson. Exploiting uniqueness in query optimization. In *Proc. IEEE Conference on Data Engineering*, pages 68–79, 1994.
- [580] J. Pearl. *Heuristics: Intelligent Search Strategies for Computer Problem Solving*. Addison Wesley, 1984.
- [581] A. Pellenkoft, C. Galindo-Legaria, and M. Kersten. Complexity of transformation-based optimizers and duplicate-free generation of alternatives. Technical Report CS-R9639, CWI, 1996.
- [582] A. Pellenkoft, C. Galindo-Legaria, and M. Kersten. The complexity of transformation-based join enumeration. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 306–315, 1997.
- [583] A. Pellenkoft, C. Galindo-Legaria, and M. Kersten. Duplicate-free generation of alternatives in transformation-based optimizers. In *Proceedings of the International Conference on Database Systems for Advanced Applications (DASFAA)*, pages 117–124, 1997.
- [584] M. Pettersson. Linux x86 performance monitoring counters driver. perform internet search for this or similar tools.
- [585] M. Pezarro. A note on estimating hit ratios for direct-access storage devices. *The Computer Journal*, 19(3):271–272, 1976.
- [586] H. Pirahesh, J. Hellerstein, and W. Hasan. Extensible/rule-based query rewrite optimization in Starburst. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 39–48, 1992.
- [587] H. Pirahesh, T. Leung, and W. Hassan. A rule engine for query transformation in Starburst and IBM DB2 C/S DBMS. In *Proc. IEEE Conference on Data Engineering*, pages 391–400, 1997.
- [588] A. Pirotte. Fundamental and secondary issues in the design of non-procedural relational languages. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 239–250, 1979.
- [589] M. Piwowarski. Comments on batched searching of sequential and tree-structured files. *ACM Trans. on Database Systems*, 10(2):285–287, 1985.
- [590] N. Plyzotis and M. Garofalakis. XSKETCH synopsis for XML. In *Hellenic Data Management Symposium 02*, 2002.
- [591] S. L. Pollack. Conversion of limited entry decision tables to computer programs. *Communications of the ACM*, 8(11):677–682, 1965.
- [592] N. Polyzotis and M. Garofalakis. Statistical synopses for graph-structured XML databases. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 358–369, 2002.

- [593] N. Polyzotis and M. Garofalakis. Structure and value synopsis for XML data graphs. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 466–477, 2002.
- [594] N. Polyzotis, M. Garofalakis, and Y. Ioannidis. Selectivity estimation for XML twigs. In *Proc. IEEE Conference on Data Engineering*, pages 264–275, 2002.
- [595] V. Poosola, Y. Ioannidis, P. Haas, and E. Shekita. Improved histograms for selectivity estimates of range predicates. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 294–305, 1996.
- [596] S. Pramanik and D. Ittner. Use of graph-theoretic models for optimal relational database accesses to perform joins. *ACM Trans. on Database Systems*, 10(1):57–74, 1985.
- [597] W. Press, S. Teukolsky, and W. Vetterling. *Numerical Recipes in C*. Cambridge University Press, 1993.
- [598] X. Qian. Query folding. In *Proc. IEEE Conference on Data Engineering*, pages 48–55, 1996.
- [599] D. Quass and J. Widom. On-line warehouse view maintenance. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 393–404, 1997.
- [600] Y.-J. Qyang. A tight upper bound for the lumped disk seek time for the SCAN disk scheduling policy. *Information Processing Letters*, 54:355–358, 1995.
- [601] E. Rahm. *Mehrrechner-Datenbanksysteme: Grundlagen der verteilten und parallelen Datenbankverwaltung*. Addison-Wesley, 1994.
- [602] A. Rajaraman, Y. Sagiv, and J.D. Ullman. Answering queries using templates with binding patterns. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, PODS, 1995.
- [603] Bernhard Mitschang Ralf Rantzau, Leonard D. Shapiro and Quan Wang. Algorithms and applications for universal quantification in relational databases. *Information Systems*, 28(1-2):3–32, 2003.
- [604] R. Ramakrishnan and J. Gehrke. *Database Management Systems*. McGraw Hill, 2000. 2nd Edition.
- [605] K. Ramamohanarao, J. Lloyd, and J. Thom. Partial-match retrieval using hashing descriptors. *ACM Trans. on Database Systems*, 8(4):552–576, 1983.
- [606] M. Ramanath, L. Zhang, J. Freire, and J. Haritsa. IMAX: Incremental maintenance of schema-based xXML statistics. In *Proc. IEEE Conference on Data Engineering*, pages 273–284, 2005.

- [607] K. Ramasamy, J. Naughton, and D. Maier. High performance implementation techniques for set-valued attributes. Technical report, University of Wisconsin, Wisconsin, 2000.
- [608] K. Ramasamy, J. Patel, J. Naughton, and R. Kaushik. Set containment joins: The good, the bad, and the ugly. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 351–362, 2000.
- [609] S. Ramaswamy and P. Kanellakis. OODB indexing by class division. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 139–150, 1995.
- [610] R. Rantzau, L. Shapiro, B. Mitschang, and Q. Wang. Universal quantification in relational databases: A classification of data and algorithms. In *Proc. of the Int. Conf. on Extending Database Technology (EDBT)*, pages 445–463, 2002.
- [611] J. Rao, B. Lindsay, G. Lohman, H. Pirahesh, and D. Simmen. Using EELs: A practical approach to outerjoin and antijoin reordering. In *Proc. IEEE Conference on Data Engineering*, pages 595–606, 2001.
- [612] J. Rao and K. Ross. Reusing invariants: A new strategy for correlated queries. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 37–48, Seattle, WA, 1998.
- [613] S. Rao, A. Badia, and D. Van Gucht. Providing better support for a class of decision support queries. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 217–227, 1996.
- [614] G. Ray, J. Haritsa, and S. Seshadri. Database compression: A performance enhancement tool. In *COMAD*, 1995.
- [615] D. Reiner and A. Rosenthal. Strategy spaces and abstract target machines for query optimization. *Database Engineering*, 5(3):56–60, Sept. 1982.
- [616] D. Reiner and A. Rosenthal. Querying relational views of networks. In W. Kim, D. Reiner, and D. Batory, editors, *Query Processing in Database Systems*, pages 109–124, 1985.
- [617] E. Reingold, J. Nievergelt, and N. Deo. *Combinatorial Algorithms: Theory and Practice*. Prentice Hall, 1977.
- [618] L. T. Reinwald and R. M. Soland. Conversion of limited entry decision tables to optimal computer programs I: minimum average processing time. *Journal of the ACM*, 13(3):339–358, 1966.
- [619] F. Reiss and T. Kanungo. A characterization of the sensitivity of query optimization to storage access cost parameters. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 385–396, 2003.

- [620] A. Reiter, A. Clute, and J. Tenenbaum. Representation and execution of searches over large tree-structured data bases. In *Proc. IFIP Congress, Booklet TA-3*, pages 134–144, 1971.
- [621] P. Richard. Evaluation of the size of a query expressed in relational algebra. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 155–163, 1981.
- [622] R. Van De Riet, A. Wassermann, M. Kersten, and W. De Jonge. High-level programming features for improving the efficiency of a relational database system. *ACM Trans. on Database Systems*, 6(3):464–485, 1981.
- [623] D.J. Rosenkrantz and M.B. Hunt. Processing conjunctive predicates and queries. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 64–74, 1980.
- [624] A. Rosenthal. Note on the expected size of a join. *SIGMOD Record*, 11(4):19–25, 1981.
- [625] A. Rosenthal and U. S. Chakravarthy. Anatomy of a modular multiple query optimizer. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 230–239, 1988.
- [626] A. Rosenthal and C. Galindo-Legaria. Query graphs, implementing trees, and freely-reorderable outerjoins. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 291–299, 1990.
- [627] A. Rosenthal, S. Heiler, U. Dayal, and F. Manola. Traversal recursion: a practical approach to supporting recursive applications. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 166–167, 1986.
- [628] A. Rosenthal and P. Helman. Understanding and extending transformation-based optimizers. *IEEE Data Engineering*, 9(4):44–51, 1986.
- [629] A. Rosenthal and D. Reiner. An architecture for query optimization. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 246–255, 1982.
- [630] A. Rosenthal and D. Reiner. Extending the algebraic framework of query processing to handle outerjoins. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 334–343, 1984.
- [631] A. Rosenthal and D. Reiner. Querying relational views of networks. In W. Kim, D. Reiner, and D. Batory, editors, *Query Processing in Database Systems*, New York, 1984. Springer.
- [632] A. Rosenthal, C. Rich, and M. Scholl. Reducing duplicate work in relational join(s): a modular approach using nested relations. Technical report, ETH Zürich, 1991.

- [633] K. Ross. Conjunctive selection conditions in main memory. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 108–120, 2002.
- [634] K. Ross, D. Srivastava, and D. Chatziantoniou. Complex aggregation at multiple granularities. In *Proc. of the Int. Conf. on Extending Database Technology (EDBT)*, pages 263–278, 1998.
- [635] M. Roth and S. Horn. Database compression. *SIGMOD Record*, 22(3):31–39, 1993.
- [636] M. Roth, H. Korth, and A. Silberschatz. Extended algebra and calculus for nested relational databases. *ACM Trans. on Database Systems*, 13(4):389–417, 1988. see also [754].
- [637] M. Roth, H. Korth, and A. Silberschatz. Null values in nested relational databases. *Acta Informatica*, 26(7):615–642, 1989.
- [638] M. Roth, H. Korth, and A. Silberschatz. Addendum to null values in nested relational databases. *Acta Informatica*, 28(6):607–610, 1991.
- [639] N. Roussopoulos. View indexing in relational databases. *ACM Trans. on Database Systems*, 7(2):258–290, 1982.
- [640] C. Ruemmler and J. Wilkes. An introduction to disk drive modeling. *IEEE Computer*, 27(3):17–29, 1994.
- [641] K. Runapongsa, J. Patel, H. Jagadish, and S. AlKhalifa. The michigan benchmark. Technical report, University of Michigan, 2002.
- [642] G. Sacco. Index access with a finite buffer. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 301–309, 1987.
- [643] G. Sacco and M. Schkolnick. A technique for managing the buffer pool in a relational system using the hot set model. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 257–262, 1982.
- [644] G. Sacco and M. Schkolnick. Buffer management in relational database systems. *ACM Trans. on Database Systems*, 11(4):473–498, 1986.
- [645] G. M. Sacco. Fragmentation: A technique for efficient query processing. *ACM Trans. on Database Systems*, 11(2):?–?, June 1986.
- [646] Y. Sagiv. *Optimization of queries in relational databases*. PhD thesis, Princeton University, 1978.
- [647] Y. Sagiv. *Optimization of Queries in Relational Databases*. UMI Research Press, Ann Arbor, Michigan, 1981.
- [648] Y. Sagiv. Quadratic algorithms for minimizing joins in restricted relational expressions. *SIAM J. Comput.*, 12(2):321–346, 1983.
- [649] Y. Sagiv and M. Yannakakis. Equivalence among expressions with the union and difference operators. *Journal of the ACM*, 27(4):633–655, 1980.

- [650] V. Sarathy, L. Saxton, and D. Van Gucht. Algebraic foundation and optimization for object based query languages. In *Proc. IEEE Conference on Data Engineering*, pages 113–133, 1993.
- [651] C. Sartiani. A general framework for estimating XML query cardinality. In *Int. Workshop on Database Programming Languages*, pages 257–277, 2003.
- [652] F. Scarcello, G. Greco, and N. Leone. Weighted hypertree decomposition and optimal query plans. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages 210–221, 2004.
- [653] J. Scheible. A survey of storage options. *IEEE Computer*, 35(12):42–46, 2002.
- [654] H.-J. Schek and M. Scholl. The relational model with relation-valued attributes. *Information Systems*, 11(2):137–147, 1986.
- [655] W. Scheufele. *Algebraic Query Optimization in Database Systems*. PhD thesis, Universität Mannheim, 1999.
- [656] W. Scheufele and G. Moerkotte. Optimal ordering of selections and joins in acyclic queries with expensive predicates. Technical Report 96-3, RWTH-Aachen, 1996.
- [657] W. Scheufele and G. Moerkotte. On the complexity of generating optimal plans with cross products. In *Proc. ACM SIGMOD/SIGACT Conf. on Princ. of Database Syst. (PODS)*, pages 238–248, 1997.
- [658] W. Scheufele and G. Moerkotte. Efficient dynamic programming algorithms for ordering expensive joins and selections. In *Proc. of the Int. Conf. on Extending Database Technology (EDBT)*, pages 201–215, 1998.
- [659] J. Schindler, A. Ailamaki, and G. Ganger. Lachesis: Robust database storage management based on device-specific performance characteristics. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 706–717, 2003.
- [660] J. Schindler and G. Ganger. Automated disk drive characterization. Technical Report CMU-CS-99-176, Carnegie Mellon University, 1999.
- [661] J. Schindler, J. Griffin, C. Lumb, and G. Ganger. Track-aligned extents: Matching access patterns to disk drive characteristics. Technical Report CMU-CS-01-119, Carnegie Mellon University, 2001.
- [662] J. Schindler, J. Griffin, C. Lumb, and G. Ganger. Track-aligned extents: Matching access patterns to disk drive characteristics. In *Conf. on File and Storage Technology (FAST)*, pages 259–274, 2002.
- [663] A. Schmidt, M. Kersten, M. Windhouwer, and F. Waas. Efficient relational storage and retrieval of XML documents. In *ACM SIGMOD Workshop on the Web and Databases (WebDB)*, 2000.

- [664] A. Schmidt, F. Waas, M. Kersten, D. Florescu, I. Manolescu, M. Carey, and R. Busse. The XML Benchmark Project. Technical Report INS-R0103, CWI, Amsterdam, 2001.
- [665] J. W. Schmidt. Some high level language constructs for data of type relation. *ACM Trans. on Database Systems*, 2(3):247–261, 1977.
- [666] M. Scholl. Theoretical foundation of algebraic optimization utilizing unnormalized relations. In *Proc. Int. Conf. on Database Theory (ICDT)*, pages ?–?, 1986.
- [667] T. Schwentick. XPath query containment. *ACM SIGMOD Record*, 33(1):101–109, 2004.
- [668] E. Sciore and J. Sieg. A modular query optimizer generator. In *Proc. IEEE Conference on Data Engineering*, pages 146–153, 1990.
- [669] B. Seeger. An analysis of schedules for performing multi-page requests. *Information Systems*, 21(4):387–407, 1996.
- [670] B. Seeger, P.-A. Larson, and R. McFadyen. Reading a set of disk pages. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 592–603, 1993.
- [671] A. Segev. Optimization of join operations in horizontally partitioned database systems. *ACM Trans. on Database Systems*, 11(1):48–80, 1986.
- [672] P. Selinger, M. Astrahan, D. Chamberlin, R. Lorie, and T. Price. Access path selection in a relational database management system. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 23–34, 1979.
- [673] T. Sellis. Intelligent caching and indexing techniques for relational database systems. *Information Systems*, 13(2):175–185, 1988.
- [674] T. Sellis. Intelligent caching and indexing techniques for relational database systems. *Information Systems*, 13(2):175–186, 1988.
- [675] T. Sellis. Multiple-query optimization. *ACM Trans. on Database Systems*, 13(1):23–52, 1988.
- [676] T. K. Sellis. Global query optimization. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 191–205, 1986.
- [677] M. Seltzer, P. Chen, and J. Ousterhout. Disk scheduling revisited. In *USENIX*, pages 313–323, 1990.
- [678] V. Sengar and J. Haritsa. PLASTIC: Reducing query optimization overheads through plan recycling. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, page 676, 2003.
- [679] K. Seppi, J. Barnes, and C. Morris. A bayesian approach to database query optimization. *ORSA Journal of Computing*, 5(4):410–418, 1993.

- [680] P. Seshadri, J. Hellerstein, H. Pirahesh, T. Leung, R. Ramakrishnan, D. Srivastava, P. Stuckey, and S. Sudarshan. Cost-based optimization for magic: Algebra and implementation. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 435–446, 1996.
- [681] P. Seshadri, H. Pirahesh, and C. Leung. Decorrelation of complex queries. Research Report RJ 9846 (86093), IBM Almaden Research Division, June 1994.
- [682] P. Seshadri, H. Pirahesh, and T. Leung. Complex query decorrelation. In *Proc. IEEE Conference on Data Engineering*, pages 450–458, 1996.
- [683] K. Sevcik. Data base system performance prediction using an analytical model. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 182–198, 1981.
- [684] D. Severance. A practitioner’s guide to data base compression. *Information Systems*, 8(1):51–62, 1983.
- [685] D. Severance and G. Lohman. Differential files: their application to the maintenance of large databases. *ACM Trans. on Database Systems*, 1(3):256–267, Sep 1976.
- [686] M. C. Shan. Optimal plan search in a rule-based query optimizer. In J. W. Schmidt, S. Ceri, and M. Missikoff, editors, *Proc. of the Intl. Conf. on Extending Database Technology*, pages 92–112, Venice, Italy, Mar 1988. Springer-Verlag, Lecture Notes in Computer Science No. 303.
- [687] J. Shanmugasundaram, R. Barr E. J. Shekita, M. J. Carey, B. G. Lindsay, H. Pirahesh, and B. Reinwald. Efficiently Publishing Relational Data as XML Documents. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 65–76, 2000.
- [688] L. Shapiro, D. Maier, P. Benninghoff, K. Billings, Y. Fan, K. Hatwal, Q. Wang, Y. Zhang, H.-M. Wu, and B. Vance. Exploiting upper and lower bounds in top-down query optimization. In *IDEAS*, pages 20–33, 2001.
- [689] L. Shapiro and A. Stephens. Bootstrap percolation, the schröder numbers and the n -kings problem. *SIAM J. Discr. Math.*, 4(2):275–280, 1991.
- [690] A. Shatdal and J. Naughton. Processing aggregates in parallel database systems. Technical Report TR 1233, University of Wisconsin, 1994.
- [691] G. M. Shaw and S.B. Zdonik. Object-oriented queries: Equivalence and optimization. In *1st Int. Conf. on Deductive and Object-Oriented Databases*, pages 264–278, 1989.
- [692] G. M. Shaw and S.B. Zdonik. A query algebra for object-oriented databases. Tech. report no. cs-89-19, Department of Computer Science, Brown University, 1989.

- [693] G.M. Shaw and S.B. Zdonik. An object-oriented query algebra. In *2nd Int. Workshop on Database Programming Languages*, pages 111–119, 1989.
- [694] G.M. Shaw and S.B. Zdonik. A query algebra for object-oriented databases. In *Proc. IEEE Conference on Data Engineering*, pages 154–162, 1990.
- [695] E. Shekita and M. Carey. A performance evaluation of pointer-based joins. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 300–311, 1990.
- [696] E. Shekita, H. Young, and K.-L. Tan. Multi-join optimization for symmetric multiprocessors. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 479–492, 1993.
- [697] P. Shenoy and H. Cello. A disk scheduling framework for next generation operating systems. In *Proc. ACM SIGMETRICS Conf. on Measurement and Modeling of Computer Systems*, pages 44–55, 1998.
- [698] S. T. Shenoy and Z. M. Ozsoyoglu. A system for semantic query optimization. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 181–195, 1987.
- [699] S. Sherman and R. Brice. Performance of a database manager in a virtual memory system. *ACM Trans. on Database Systems*, 1(4):317–343, 1976.
- [700] B. Shneiderman and V. Goodman. Batched searching of sequential and tree structured files. *ACM Trans. on Database Systems*, 1(3):208–222, 1976.
- [701] E. Shriver. *Performance Modeling for Realistic Storage Devices*. PhD thesis, University of New York, 1997.
- [702] E. Shriver, A. Merchant, and J. Wilkes. An analytical behavior model for disk drives with readahead caches and request reordering. In *Proc. ACM SIGMETRICS Conf. on Measurement and Modeling of Computer Systems*, pages 182–191, 1998.
- [703] A. Shrufi and T. Topaloglou. Query processing for knowledge bases using join indices. In *Int. Conference on Information and Knowledge Management (CIKM)*, 1995.
- [704] A. Shukla, P. Deshpande, J. Naughton, and K. Ramasamy. Storage estimation for multidimensional aggregates in the presence of hierarchies. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, 1996.
- [705] K. Shwayder. Conversion of limited entry decision tables to computer programs — a proposed modification to Pollack’s algorithm. *Communications of the ACM*, 14(2):69–73, 1971.
- [706] M. Siegel, E. Sciore, and S. Salveter. A method for automatic rule derivation to support semantic query optimization. *ACM Trans. on Database Systems*, 17(4):53–600, 1992.

- [707] A. Silberschatz, H. Korth, and S. Sudarshan. *Database System Concepts*. McGraw Hill, 1997. 3rd Edition.
- [708] D. Simmen, C. Leung, and H. Pirahesh. Exploitation of uniqueness properties for the optimization of SQL queries using a 1-tuple condition. Research Report RJ 10008 (89098), IBM Almaden Research Division, Feb. 1996.
- [709] D. Simmen, E. Shekita, and T. Malkemus. Fundamental techniques for order optimization. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 57–67, 1996.
- [710] D. Simmen, E. Shekita, and T. Malkemus. Fundamental techniques for order optimization. In *Proc. of the Int. Conf. on Extending Database Technology (EDBT)*, pages 625–62, 1996.
- [711] G. Slivinskas, C. Jensen, and R. Snodgrass. Bringing order to query optimization. *SIGMOD Record*, 13(2):5–14, 2002.
- [712] D. Smith and M. Genesereth. Ordering conjunctive queries. *Artificial Intelligence*, 26:171–215, 1985.
- [713] J. A. Smith. Sequentiality and prefetching in database systems. *ACM Trans. on Database Systems*, 3(3):223–247, 1978.
- [714] J. M. Smith and P. Y.-T. Chang. Optimizing the performance of a relational algebra database interface. *Communications of the ACM*, 18(10):568–579, 1975.
- [715] R. Sasic, J. Gu, and R. Johnson. The Unison algorithm: Fast evaluation of boolean expressions. *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, 1:456 – 477, 1996.
- [716] D. Srivastava, S. Al-Khalifa, H. V. Jagadish, N. Koudas, J. Patel, and Y. Wu. Structural joins: A primitive for efficient XML query pattern matching. In *Proc. IEEE Conference on Data Engineering*, 2002.
- [717] D. Srivastava, S. Dar, J. Jagadish, and A. Levy. Answering queries with aggregation using views. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 318–329, 1996.
- [718] R. Stanley. *Enumerative Combinatorics, Volume I*, volume 49 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, 1997.
- [719] H. Steenhagen. *Optimization of Object Query Languages*. PhD thesis, University of Twente, 1995.
- [720] H. Steenhagen, P. Apers, and H. Blanken. Optimization of nested queries in a complex object model. In *Proc. of the Int. Conf. on Extending Database Technology (EDBT)*, pages 337–350, 1994.

- [721] H. Steenhagen, P. Apers, H. Blanken, and R. de By. From nested-loop to join queries in oodb. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 618–629, 1994.
- [722] H. Steenhagen, R. de By, and H. Blanken. Translating OSQL queries into efficient set expressions. In *Proc. of the Int. Conf. on Extending Database Technology (EDBT)*, pages 183–197, 1996.
- [723] M. Steinbrunn, G. Moerkotte, and A. Kemper. Heuristic and randomized optimization for the join ordering problem. *The VLDB Journal*, 6(3):191–208, Aug. 1997.
- [724] M. Steinbrunn, K. Peithner, G. Moerkotte, and A. Kemper. Bypassing joins in disjunctive queries. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 228–238, 1995.
- [725] K. Stocker, D. Kossmann, R. Braumandl, and A. Kemper. Integrating semi-join reducers into state-of-the-art query processors. In *Proc. IEEE Conference on Data Engineering*, pages 575–584, 2001.
- [726] L. Stockmeyer and C. Wong. On the number of comparisons to find the intersection of two relations. Technical report, IBM Watson Research Center, 1978.
- [727] H. Stone and H. Fuller. On the near-optimality of the shortest-latency-time-first drum scheduling discipline. *Communications of the ACM*, 16(6):352–353, 1973.
- [728] M. Stonebraker. Inclusion of new types in relational database systems. In *Proc. IEEE Conference on Data Engineering*, pages ?–?, 1986.
- [729] M. Stonebraker, J. Anton, and E. Hanson. Extending a database system with procedures. *ACM Trans. on Database Systems*, 12(3):350–376, Sep 1987.
- [730] M. Stonebraker and P. Brown. *Object-Relational DBMSs, Tracking the Next Great Wave*. Morgan Kaufman, 1999.
- [731] M. Stonebraker et al. QUEL as a data type. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, Boston, MA, June 1984.
- [732] M. Stonebraker, A. Jhingran, J. Goh, and S. Potamios. On rules, procedures, caching and views in data base systems. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 281–290, 1990.
- [733] M. Stonebraker and L. A. Rowe. The design of postgres. In *Proc. of the 15th ACM SIGMOD*, pages 340–355, 1986.
- [734] M. Stonebraker, E. Wong, P. Kreps, and G. Held. The design and implementation of INGRES. *ACM Trans. on Database Systems*, 1(3):189–222, 1976.

- [735] D. Straube and T. Özsu. Access plan generation for an object algebra. Technical Report TR 90-20, Department of Computing Science, University of Alberta, June 1990.
- [736] D. Straube and T. Özsu. Queries and query processing in object-oriented database systems. Technical report, Department of Computing Science, University of Alberta, Edmonton, Alberta, Canada, 1990.
- [737] D. Straube and T. Özsu. Queries and query processing in object-oriented database systems. *ACM Trans. on Information Systems*, 8(4):387–430, 1990.
- [738] D. Straube and T. Özsu. Execution plan generation for an object-oriented data model. In *Proc. Int. Conf. on Deductive and Object-Oriented Databases (DOOD)*, pages 43–67, 1991.
- [739] D. D. Straube. *Queries and Query Processing in Object-Oriented Database Systems*. PhD thesis, The University of Alberta, Edmonton, Alberta, Canada, Dec 1990.
- [740] S. Subramanian and S. Venkataraman. Cost-based optimization of decision support queries using transient views. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 319–330, Seattle, WA, 1998.
- [741] D. Suci. Query decomposition and view maintenance for query languages for unconstrained data. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 227–238, 1996.
- [742] N. Südkamp and V. Linnemann. Elimination of views and redundant variables in an SQL-like database language for extended NF² structures. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 302–313, 1990.
- [743] K. Sutner, A. Satyanarayana, and C. Suffel. The complexity of the residual node connectedness reliability problem. *SIAM J. Comp.*, 20(1):149–155, 1991.
- [744] P. Svensson. On search performance for conjunctive queries in compressed, fully transposed ordered files. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 155–163, 1979.
- [745] A. Swami. *Optimization of Large Join Queries*. PhD thesis, Stanford University, 1989. Technical Report STAN-CS-89-1262.
- [746] A. Swami. Optimization of large join queries: Combining heuristics and combinatorial techniques. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 367–376, 1989.
- [747] A. Swami and A. Gupta. Optimization of large join queries. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 8–17, 1988.

- [748] A. Swami and B. Iyer. A polynomial time algorithm for optimizing join queries. Technical Report RJ 8812, IBM Almaden Research Center, 1992.
- [749] A. Swami and B. Iyer. A polynomial time algorithm for optimizing join queries. In *Proc. IEEE Conference on Data Engineering*, pages 345–354, 1993.
- [750] A. Swami and B. Schiefer. Estimating page fetches for index scans with finite LRU buffers. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 173–184, 1994.
- [751] A. Swami and B. Schiefer. On the estimation of join result sizes. In *Proc. of the Int. Conf. on Extending Database Technology (EDBT)*, pages 287–300, 1994.
- [752] N. Talagala, R. Arpaci-Dusseau, and D. Patterson. Microbenchmark-based extraction of local and global disk characteristics. Technical Report UCB-CSD-99-1063, University of Berkeley, 2000.
- [753] K.-L. Tan and H. Lu. A note on the strategy space of multiway join query optimization problem in parallel systems. *SIGMOD Record*, 20(4):81–82, 1991.
- [754] A. Tansel and L. Garnett. On roth, korth, and silberschatz’s extended algebra and calculus for nested relational databases. *ACM Trans. on Database Systems*, 17(2):374–383, 1992.
- [755] Y. C. Tay. On the optimality of strategies for multiple joins. *Journal of the ACM*, 40(5):1067–1086, 1993.
- [756] T. Teorey and T. Pinkerton. A comparative analysis of disk scheduling policies. In *Proc. of the AFIPS Fall Joint Computer Conference*, pages 1–11, 1972.
- [757] T. Teorey and T. Pinkerton. A comparative analysis of disk scheduling policies. *Communications of the ACM*, 15(3):177–184, 1972.
- [758] J. Teubner, T. Grust, and M. Van Keulen. Bridging the gap between relational and native XML storage with staircase join. *Grundlagen von Datenbanken*, pages 85–89, 2003.
- [759] C. Tompkins. Machine attacks o problems whose variables are permutations. *Numerical Analysis (Proc. of Symposia in Applied Mathematics)*, 6, 1956.
- [760] R. Topor. Join-ordering is NP-complete. Draft, personal communication, 1998.
- [761] Transaction Processing Council (TPC). TPC Benchmark D. <http://www.tpc.org>, 1995.

- [762] Transaction Processing Performance Council, 777 N. First Street, Suite 600, San Jose, CA, USA. *TPC Benchmark R*, 1999. Revision 1.2.0. <http://www.tpc.org>.
- [763] P. Triantafillou, S. Christodoulakis, and C. Georgiadis. A comprehensive analytical performance model for disk devices under random workloads. *IEEE Trans. on Knowledge and Data Eng.*, 14(1):140–155, 2002.
- [764] O. Tsatalos, M. Solomon, and Y. Ioannidis. The GMAP: A versatile tool for physical data independence. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 367–378, 1994.
- [765] K. Tufte and D. Maier. Aggregation and accumulation of XML data. *IEEE Data Engineering Bulletin*, 24(2):34–39, 2001.
- [766] J.D. Ullman. *Database and Knowledge Base Systems*, volume Volume 1. Computer Science Press, 1989.
- [767] J.D. Ullman. *Database and Knowledge Base Systems*, volume Volume 2. Computer Science Press, 1989.
- [768] J.D. Ullman. *Database and Knowledge Base Systems*. Computer Science Press, 1989.
- [769] D. Straube und T. Özsu. Query transformation rules for an object algebra. Technical Report TR 89-23, Department of Computing Science, University of Alberta, Sept. 1989.
- [770] T. Urhan, M. Franklin, and L. Amsaleg. Cost based query scrambling for initial delays. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 130–141, 1998.
- [771] M. Uysal, G. Alvarez, and A. Merchant. A modular analytical throughput model for modern disk arrays. In *MASCOTS*, pages 183–192, 2001.
- [772] P. Valduriez. Join indices. *ACM Trans. on Database Systems*, 12(2):218–246, 1987.
- [773] P. Valduriez and H. Boral. Evaluation of recursive queries using join indices. In *Proc. Int. Conf. on Expert Database Systems (EDS)*, pages 197–208, 1986.
- [774] P. Valduriez and S. Danforth. Query optimization in database programming languages. In *Proc. Int. Conf. on Deductive and Object-Oriented Databases (DOOD)*, pages 516–534, 1989.
- [775] L. Valiant. The complexity of computing the permanent. *Theoretical Comp. Science*, 8:189–201, 1979.
- [776] L. Valiant. The complexity of enumeration and reliability problems. *SIAM J. Comp.*, 8(3):410–421, 1979.

- [777] B. Vance. *Join-order Optimization with Cartesian Products*. PhD thesis, Oregon Graduate Institute of Science and Technology, 1998.
- [778] B. Vance and D. Maier. Rapid bushy join-order optimization with cartesian products. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 35–46, 1996.
- [779] S. L. Vandenberg and D. DeWitt. An algebra for complex objects with arrays and identity. Internal report, Computer Sciences Department, University of Wisconsin, Madison, WI 53706, USA, 1990.
- [780] S. L. Vandenberg and D. DeWitt. Algebraic support for complex objects with arrays, identity, and inheritance. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 158–167, 1991.
- [781] G. von Bülzingsloewen. Translating and optimizing sql queries having aggregates. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 235–243, 1987.
- [782] G. von Bülzingsloewen. *Optimierung von SQL-Anfragen für parallele Bearbeitung (Optimization of SQL-queries for parallel processing)*. PhD thesis, University of Karlsruhe, 1990. in German.
- [783] G. von Bülzingsloewen. *SQL-Anfragen: Optimierung für parallele Bearbeitung*. FZI-Berichte Informatik. Springer, 1991.
- [784] F. Waas and A. Pellenkoft. Probabilistic bottom-up join order selection – breaking the curse of NP-completeness. Technical Report INS-R9906, CWI, 1999.
- [785] F. Waas and A. Pellenkoft. Join order selection - good enough is easy. In *BNCOD*, pages 51–67, 2000.
- [786] J. Wang, J. Li, and G. Butler. Implementing the PostgreSQL query optimizer within the OPT++ framework. In *Asia-Pacific Software Engineering Conference (APSEC)*, pages 262–272, 2003.
- [787] J. Wang, M. Maher, and R. Topor. Rewriting unions of general conjunctive queries using views. In *Proc. of the Int. Conf. on Extending Database Technology (EDBT)*, pages 52–69, 2002.
- [788] W. Wang, H. Jiang, H. Lu, and J. Yu. Containment join size estimation: Models and methods. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 145–156, 2003.
- [789] W. Wang, H. Jiang, H. Lu, and J. Yu. Bloom histogram: Path selectivity estimation for xml data with updates. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 240–251, 2004.
- [790] X. Wang and M. Cherniack. Avoiding ordering and grouping in query processing. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 826–837, 2003.

- [791] S. Waters. File design fallacies. *The Computer Journal*, 15(1):1–4, 1972.
- [792] S. Waters. Hit ratio. *Computer Journal*, 19(1):21–24, 1976.
- [793] H. Wedekind and G. Zörntlein. Prefetching in realtime database applications. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 215–226, 1986.
- [794] M. Wedekind. On the selection of access paths in a database system. In J. Klimbie and K. Koffeman, editors, *IFIP Working Conference Data Base Management*, pages 385–397, Amsterdam, 1974. North-Holland.
- [795] G. Weikum. Set-oriented disk access to large complex objects. In *Proc. IEEE Conference on Data Engineering*, pages 426–433, 1989.
- [796] G. Weikum, B. Neumann, and H.-B. Paul. Konzeption und Realisierung einer mengenorientierten Seitenschnittstelle zum effizienten Zugriff auf komplexe Objekte. In *Proc. der GI-Fachtagung Datenbanksysteme für Büro, Technik und Wissenschaft (BTW)*, pages 212–230, 1987.
- [797] T. Westmann. *Effiziente Laufzeitsysteme für Datenlager*. PhD thesis, University of Mannheim, 2000.
- [798] T. Westmann, D. Kossmann, S. Helmer, and G. Moerkotte. The implementation and performance of compressed databases. Technical Report 03/98, University of Mannheim, 1998.
- [799] T. Westmann, D. Kossmann, S. Helmer, and G. Moerkotte. The implementation and performance of compressed databases. *SIGMOD Record*, 29(3):55–67, 2000.
- [800] T. Westmann and G. Moerkotte. Variations on grouping and aggregations. Technical Report 11/99, University of Mannheim, 1999.
- [801] K.-Y. Whang, A. Malhotra, G. Sockut, and L. Burns. Supporting universal quantification in a two-dimensional database query language. In *Proc. IEEE Conference on Data Engineering*, pages 68–75, 1990.
- [802] K.-Y. Whang, G. Wiederhold, and D. Sagalowicz. Estimating block accesses in database organizations: A closed noniterative formula. *Communications of the ACM*, 26(11):940–944, 1983.
- [803] N. Wilhelm. A general model for the performance of disk systems. *Journal of the ACM*, 24(1):14–31, 1977.
- [804] D. E. Willard. Efficient processing of relational calculus queries using range query theory. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 164–175, 1984.
- [805] C. Williams and T. Hogg. Using deep structure to locate hard problems. In *Proc. National Conference on Artificial Intelligence*, pages 472–477, 1992.

- [806] J. Wolf, R. Iyer, K. Pattipati, and J. Turek. Optimal buffer partitioning for the nested block join algorithm. In *Proc. IEEE Conference on Data Engineering*, pages 510–519, 1991.
- [807] C. Wong. Minimizing expected head movement in one-dimensional and two-dimensional mass storage systems. *ACM Computing Surveys*, 12(2):167–177, 1980.
- [808] F. Wong and K. Youssefi. Decomposition – a strategy for query processing. *ACM Trans. on Database Systems*, 1(3):223–241, 1976.
- [809] H. Wong and J. Li. Transposition algorithms on very large compressed databases. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 304–311, 1986.
- [810] P. Wood. On the equivalence of XML patterns. In *CL 2000*, 2000.
- [811] P. Wood. Minimizing simple XPath expressions. In *Int. Workshop on Database Programming Languages*, pages 13–18, 2001.
- [812] P. Wood. Containment for XPath fragments under dtd constraints. In *Proc. Int. Conf. on Database Theory (ICDT)*, pages 300–314, 2003.
- [813] W. A. Woods. Procedural semantics for question-answering systems. In *FJCC (AFIPS Vol. 33 Part I)*, pages 457–471, 1968.
- [814] B. Worthington, G. Ganger, Y. Patt, and J. Wilkes. Scheduling algorithms for modern disk drives. In *Proc. ACM SIGMETRICS Conf. on Measurement and Modeling of Computer Systems*, pages 241–251, 1994.
- [815] B. Worthington, G. Ganger, Y. Patt, and J. Wilkes. On-line extraction of SCSI disk drive parameters. In *Proc. ACM SIGMETRICS Conf. on Measurement and Modeling of Computer Systems*, pages 146–156, 1995.
- [816] B. Worthington, G. Ganger, Y. Patt, and J. Wilkes. On-line extraction of SCSI disk drive parameters. Technical Report CSE-TR-323-96, University of Michigan, 1996.
- [817] M.-C. Wu. Query optimization for selections using bitmaps. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 227–238, 1999.
- [818] Y. Wu, J. Patel, and H.V. Jagadish. Estimating answer sizes for XML queries. In *Proc. of the Int. Conf. on Extending Database Technology (EDBT)*, pages 590–608, 2002.
- [819] Y. Wu, J. Patel, and H.V. Jagadish. Estimating answer sizes for XML queries. *Information Systems*, 28(1-2):33–59, 2003.
- [820] Z. Xie. Optimization of object queries containing encapsulated methods. In *Proc. 2nd. Int. Conf. on Information and Knowledge Management*, pages 451–460, 1993.

- [821] G. D. Xu. Search control in semantic query optimization. Technical Report 83-09, COINS, University of Massachusetts, Amherst, MA, 1983.
- [822] W. Yan. *Rewriting Optimization of SQL Queries Containing GROUP-BY*. PhD thesis, University of Waterloo, 1995.
- [823] W. Yan and P.-A. Larson. Performing group-by before join. Technical Report CS 93-46, Dept. of Computer Science, University of Waterloo, Canada, 1993.
- [824] W. Yan and P.-A. Larson. Performing group-by before join. In *Proc. IEEE Conference on Data Engineering*, pages 89–100, 1994.
- [825] W. Yan and P.-A. Larson. Eager aggregation and lazy aggregation. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 345–357, 1995.
- [826] W. Yan and P.-A. Larson. Interchanging the order of grouping and join. Technical Report CS 95-09, Dept. of Computer Science, University of Waterloo, Canada, 1995.
- [827] H. Yang and P.-A. Larson. Query transformation for PSJ-queries. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 245–254, 1987.
- [828] J. Yang, K. Karlapalem, and Q. Li. Algorithms for materialized view design in data warehousing environment. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 136–145, 1997.
- [829] Qi Yang. Computation of chain queries in distributed database systems. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 348–355, 1994.
- [830] M. Yannakakis. Algorithms for acyclic database schemes. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 82–94, 1981.
- [831] B. Yao and T. Özsu. XBench – A Family of Benchmarks for XML DBMSs. Technical Report CS-2002-39, University of Waterloo, 2002.
- [832] B. Yao, T. Özsu, and N. Khandelwal. Xbench benchmark and performance testing of XML DBMSs. In *Proc. IEEE Conference on Data Engineering*, pages 621–632, 2004.
- [833] S. B. Yao. Approximating block accesses in database organizations. *Communications of the ACM*, 20(4):260–261, 1977.
- [834] S. B. Yao. An attribute based model for database access cost analysis. *ACM Trans. on Database Systems*, 2(1):45–67, 1977.
- [835] S. B. Yao and D. DeJong. Evaluation of database access paths. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 66–77, 1978.
- [836] S.B. Yao. Optimization of query evaluation algorithms. *ACM Trans. on Database Systems*, 4(2):133–155, 1979.

- [837] S.B. Yao, A.R. Hevner, and H. Young-Myers. Analysis of database system architectures using benchmarks. *IEEE Trans. on Software Eng.*, SE-13(6):709–725, 1987.
- [838] Y. Yoo and S. Lafortune. An intelligent search method for query optimization by semijoins. *IEEE Trans. on Knowledge and Data Eng.*, 1(2):226–237, June 1989.
- [839] M. Yoshikawa, T. Amagasa, T. Shimura, and S. Uemura. Xrel: A path-based approach to storage and retrieval of XML documents using relational databases. *ACM Transactions on Internet Technology*, 1(1):110–141, June 2001.
- [840] K. Youssefi and E. Wong. Query processing in a relational database management system. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 409–417, 1979.
- [841] L. Yu and S. L. Osborn. An evaluation framework for algebraic object-oriented query models. In *Proc. IEEE Conference on Data Engineering*, 1991.
- [842] J. Zahorjan, B. Bell, and K. Sevcik. Estimating block transfers when record access probabilities are non-uniform. *Information Processing Letters*, 16(5):249–252, 1983.
- [843] B. T. Vander Zander, H. M. Taylor, and D. Bitton. Estimating block accesses when attributes are correlated. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, pages 119–127, 1986.
- [844] S. Zdonik and G. Mitchell. ENCORE: An object-oriented approach to database modelling and querying. *IEEE Data Engineering Bulletin*, 14(2):53–57, June 1991.
- [845] N. Zhang, V. Kacholia, and T. Özsu. A succinct physical storage scheme for efficient evaluation of path queries in XML. In *Proc. IEEE Conference on Data Engineering*, pages 54–65, 2004.
- [846] N. Zhang and T. Özsu. Optimizing correlated path expressions in XML languages. Technical Report CS-2002-36, University of Waterloo, 2002.
- [847] Y. Zhao, P. Deshpande, and J. Naughton. An array-based algorithm for simultaneous multidimensional aggregates. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 159–170, 1997.
- [848] Y. Zhao, P. Deshpande, J. Naughton, and A. Shukla. Simultaneous optimization and evaluation of multiple dimensional queries. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, pages 271–282, 1998.
- [849] Y. Zhuge, H. Garcia-Molina, J. Hammer, and J. Widom. View maintenance in a warehouse environment. In *Proc. of the ACM SIGMOD Conf. on Management of Data*, 1995.

Appendix E

ToDo

- size of a query in rel alg: [621]
- [787]
- Integrating Buffer Issues into Query Optimization: [180, 404]
- Integrating concurrency control issues into query optimization: [531, 532]
- [86]
- where do we put "counting page accesses"?
- control, A^* , ballooning: [523, 522]
- Bypass Plans
- Properties (rather complete list, partial ordering, plan independent properties: store them somewhere else (dpstructure or memostructure))
- describe prep-phase of plan generator
- reuse plans: [678]
- estimating query compilation time: [380]
- cost model [683]
- sensitivity of QO to storage access cost parameters [619] (and join selectivities on join order: [444] [papier ist nicht ernst zu nehmen])
- magic set and semi join reducers [70, 72, 71, 149, 288, 538, 536, 538, 537, 680, 725, 838]
- join indexes and clustering tuples of different relations with 1:n relationship [197, 345, 772, 773, 703]
- B-Trees with space filling curves (Bayer)
- Prefetching [793]

- feedback to optimizer [436]
- compression [26, 48, 141, 177, 222, 221, 283, 289] [547, 614, 635, 684, 685, 744, 798, 809]
- semantic QO SQO: [1, 73, 146, 279, 311, 425, 426, 441, 447] [548, 556, 557, 579, 698, 706, 708, 821] [475]
- join processing with nonclustered indexes: [549]
- join+buffer: [806]
- removal/elimination of redundant joins [558, 742]
- benchmark(ing): Gray Book: [312]; papers: [83, 91, 566, 641, 664, 837, 831, 832]
- dynamic qo: [570] [28, 770] [38] [405]
- unnesting: [574, 612]
- prefetching: [578, 577, 713, 793]
- Starburst: [586, 587]
- BXP: [90, 223, 265, 322, 341, 376, 420, 591, 618, 705, 712, 715, 419]
BXP complexity: [68] BXP var inf: [406]
- joins: [596]
- query folding: [598]
- quantification: [94, 93, 157, 158, 610, 603, 801] [397]
- outerjoins: [75, 76, 186, 263, 255, 254, 611, 626]
- partial match + hashing: [605]
- OODB indexing by class division: [164, 609]
- decision support [613]
- tree structured databases: [620]
- Rosenthal: [629, 615, 630, 631, 616, 625, 628, 627]
- conj. queries [623]
- aggregation/(generalized proj): [97, 164, 256, 634] [326, 327, 350]
- do nest/unnest to optimize duplicate work: [632]
 $e_1 \bowtie_{A_1=A_2} e_2 \equiv \mu_g(e_1 \bowtie_{A_1=A_2} \Gamma_{g:=A_2;id}(e_2))$
- join size: [624]
- fragmentation: [645]

- eqv: [20, 19]
- alg eqvs union/difference: [649] [787]
- other sagiv: [647, 648]
- bayesian approach to QO: [679]
- cache query plans: [678]
- joins for horizontally fragmentation: [671]
- partitioning: [49, 89, 339, 409, 551]
- MQO: [25, 115, 113, 676, 675, 848]
- indexing+caching: [674]
- rule-based QO: [686, 59, 60, 251]
- rule-based IRIS: [196]
- cost: [726] [771]
- search space: [753], join ordering: [755]
- access path: [109, 794, 835, 84]
- eff aggr: [253] [800]
- misc: [804] [9] [13]
- access paths: bitmaps [817]
- dist db: [33, 34, 82, 184, 829] Donald's state of the art: [439]
- [121, 122]
- eqv: bags [22, 189]
- eqvs old: [23]
- DB2: norwegian analysis: [29]
- nested: [41]
- Genesis/Praire/Batory: [50, 54, 53, 55, 185]
- eqvs OO: [61, 62]
- dupelim: [78]
- (generalized) division: [107, 183, 301, 291]
- early aggregation
- chunks-wise processing [198, 298]

- temporal intersection join: [323]
- 2nd ord sig: Güting: [331]
- classics: [336]
- smallest first: [340]
- Hwang/Yu: [375]
- Kambayashi: [408]
- Koch [434], Lehnert [465]
- I/O cost reduction for (hash) joins: [486, 516]
- dist nest: [227]
- band join: [491]
- Donovan (TODS 76,1,4) Decision Support: [212]
- whenever materialize something (sort, hash join, etc) compute min/max of some attributes and use these as additional selection predicates
- determine optimal page access sequence and buffer size to access pairs (x,y) of pages where join partners of one relation lie on x and of the other on y (Fotouhi, Pramanik [248], Merret, Kambayashi, Yasuura [516], Omiecinski [549], Pramanik, Ittner [596], Chan, Ooi [117])
- Scarcello, Greco, Leone: [652]
- Sigmod05:
 - proactive reoptimization [40]
 - robust query optimizer [39]
 - stacked indexed views [193]
 - NF²-approach to processing nested sql queries [100]
 - efficient computatio of multiple groupby queries [143]
- LOCI: [569]
- Wavelet synopses: [274]
- Progress Indicators: [127, 493, 494]
- PostgresExperience: [786]
- Chaudhuri sigmod 05: [39]