


S.Ceri G.Gottlob L.Tanca

Logic Programming and Databases



With 42 Figures



Springer-Verlag Berlin Heidelberg New York
London Paris Tokyo Hong Kong

Table of Contents

Chapter 1

Logic Programming and Databases: An Overview	1
1.1 Logic Programming as Query Language	2
1.2 Prolog and Datalog	9
1.3 Alternative Architectures	11
1.4 Applications	14
1.5 Bibliographic Notes	14

Chapter 2

A Review of Relational Databases and Prolog	16
2.1 Overview of Relational Databases	16
2.1.1 The Relational Model	16
2.1.2 Relational Languages	18
2.2 Prolog: A Language for Programming in Logic	23
2.3 Bibliographic Notes	26

Part I Coupling Prolog to Relational

Databases	27
----------------------------	-----------

Chapter 3

Prolog as a Query Language	29
3.1 The Anti-Trust Control Problem	30
3.2 The Bill of Materials Problem	34
3.3 Conclusions	38
3.4 Bibliographic Notes	38
3.5 Exercises	38

Chapter 4

Coupling Prolog Systems to Relational Databases 40

4.1	Architectures for Coupling Prolog and Relational Systems	40
4.1.1	Assumptions and Terminology	40
4.1.2	Components of a CPR System	42
4.1.3	Architecture of CPR Systems	45
4.2	Base Conjunctions	47
4.2.1	Determining Base Conjunctions in LCPR Systems . .	49
4.2.2	Improving Base Conjunctions in TCPR Systems . . .	54
4.3	Optimization of the Prolog/Database Interface	57
4.3.1	Caching of Data	58
4.3.2	Caching of Data and Queries	58
4.3.3	Use of Subsumption	59
4.3.4	Caching Queries	60
4.3.5	Parallelism and Pre-fetching in Database Interfaces .	61
4.4	Conclusions	62
4.5	Bibliographic Notes	62
4.6	Exercises	63

Chapter 5

Overview of Systems for Coupling Prolog to Relational Databases 65

5.1	PRO-SQL	65
5.2	EDUCE	67
5.3	ESTEAM	68
5.4	BERMUDA	69
5.5	CGW and PRIMO	70
5.6	QUINTUS-PROLOG	72
5.7	Bibliographic Notes	74

Part II Foundations of Datalog 75

Chapter 6

Syntax and Semantics of Datalog 77

6.1	Basic Definitions and Assumptions	77
6.1.1	Alphabets, Terms, and Clauses	77
6.1.2	Extensional Databases and Datalog Programs	81
6.1.3	Substitutions, Subsumption, and Unification	83

6.2	The Model Theory of Datalog	80
6.2.1	Possible Worlds, Truth, and Herbrand Interpretations	80
6.2.2	The Least Herbrand Model	91
6.3	Conclusions	91
6.4	Bibliographic Notes	91
6.5	Exercises	91

Chapter 7

Proof Theory and Evaluation Paradigms of Datalog 91

7.1	The Proof Theory of Datalog	91
7.1.1	Fact Inference	91
7.1.2	Soundness and Completeness of the Inference Rule EP	91
7.2	Least Fixpoint Iteration	101
7.2.1	Basic Results of Fixpoint Theory	101
7.2.2	Least Fixpoints and Datalog Programs	101
7.3	Backward Chaining and Resolution	101
7.3.1	The Principle of Backward Chaining	101
7.3.2	Resolution	111
7.4	Conclusions	121
7.5	Bibliographic Notes	121
7.6	Exercises	121

Part III Optimization Methods for Datalog 121

Chapter 8

Classification of Optimization Methods for Datalog 121

8.1	Criteria for the Classification of Optimization Methods	121
8.1.1	Formalism	121
8.1.2	Search Strategy	121
8.1.3	Objectives of Optimization Methods	121
8.1.4	Type of Information Considered	121
8.2	Classification of Optimization Methods	121
8.3	Translation of Datalog into Relational Algebra	130
8.4	Classification of Datalog Rules	136
8.5	The Expressive Power of Datalog	142
8.6	Bibliographic Notes	143
8.7	Exercises	144

Chapter 9

Evaluation Methods	145
9.1 Bottom-up Evaluation	145
9.1.1 Algebraic Naive Evaluation	145
9.1.2 Semi-naive Evaluation	150
9.1.3 The Method of Henschen and Naqvi	154
9.2 Top-down Evaluation	155
9.2.1 Query-Subquery	155
9.2.2 The RQA/FQI Method	160
9.3 Bibliographic Notes	161
9.4 Exercises	162

Chapter 10

Rewriting Methods	163
10.1 Logical Rewriting Methods	163
10.1.1 Magic Sets	165
10.1.2 The Counting Method	174
10.1.3 The Static Filtering Method	177
10.1.4 Semi-naive Evaluation by Rewriting	183
10.2 Rewriting of Algebraic Systems	185
10.2.1 Reduction to Union-Join Normal Form	185
10.2.2 Determination of Common Subexpressions	187
10.2.3 Query Subsetting and Strong Components	189
10.2.4 Marking of Variables	191
10.2.5 Reduction of Variables	193
10.2.6 Reduction of Constants	193
10.2.7 Summary of the Algebraic Approach	200
10.3 A General View of Optimization	200
10.4 Bibliographic Notes	205
10.5 Exercises	206

Chapter 11

Extensions of Pure Datalog	208
11.1 Using Built-in Predicates in Datalog	210
11.2 Incorporating Negation into Datalog	211
11.2.1 Negation and the Closed World Assumption	212
11.2.2 Stratified Datalog	215
11.2.3 Perfect Models and Local Stratification	224
11.2.4 Inflationary Semantics and Expressive Power	226

11.3	Representation and Manipulation of Complex Objects	228
11.3.1	Basic Features of LDL	229
11.3.2	Semantics of Admissible LDL Programs	235
11.3.3	Data Models for Complex Objects	240
11.4	Conclusions	241
11.5	Bibliographic Notes	241
11.6	Exercises	244

Chapter 12

Overview of Research Prototypes for Integrating Relational Databases and Logic Programming	246
---	------------

12.1	The LDL Project	247
12.2	The NAIL! Project	251
12.3	The POSTGRES Project	255
12.4	The FIFTH GENERATION Project	257
12.5	The KIWI Project	260
12.6	The ALGRES Project	262
12.7	The PRISMA Project	264
12.8	Bibliographic Notes	265

Bibliography	267
-------------------------------	------------

Index	277
------------------------	------------