

CONTENTS

Session 1: Temporal Databases

The Time Index - An Access Structure for Temporal Data	1
R Elmasri, G T Wu and Y G Kim	
A Temporal Relational Algebra as a Basis for Temporal Relational Completeness .	13
A Tuzhilin and J Clifford	
Support for Temporal Data by Complex Objects	24
W Käfer, N Ritter and H Schöning	

Session 2: Indexing Techniques

Indexing in a Hypertext Database	36
C Clifton and H Garcia-Molina	
Search Key Substitution in the Encipherment of B-Trees	50
T Hardjono and J Seberry	
Query Processing Method for Multi-Attribute Clustered Relations	59
L Harada, M Nakano, M Kitsuregawa and M Takagi	

Session 3: Transaction Processing 1

Triggered Real-Time Databases with Consistency Constraints	71
H F Korth, N Soparkar and A Silberschatz	
Cooperative Transaction Hierarchies: A Transaction Model to Support Design Applications	83
M H Nodine and S B Zdonik	
A Formal Approach to Recovery by Compensating Transactions	95
H F Korth, E Levy and A Silberschatz	

Session 4: Performance of OODB Systems

A Study of Three Alternative Workstation Server Architectures for Object-Oriented Database Systems107
D J DeWitt, P Fattersack, D Maier and F Velez

Maintaining Consistency of Client-Cached Data 122
K Wilkinson and M-A Neimat

The Performance and Utility of the Cactis Implementation Algorithms 135
P Drew, R King and S Hudson

Session 5: I/O Architecture

Parity Striping of Disc Arrays: Low-Cost Reliable Storage with Acceptable Throughput 148
J Gray, B Horst and M Walker

Performance Analysis of Disk Arrays Under Failure162
R R Muntz and J C S Lui

Transaction Support in Read Optimized and Write Optimized File Systems 174
M Seltzer and M Stonebraker

Session 6: Hash-Based Join Algorithms

An Adaptive Hash Join Algorithm for Multiuser Environments 186
H Zeller and J Gray

Hash-Based Join Algorithms for Multiprocessor Computers with Shared Memory .198
H Lu and K L Tan and M-C Shan

Bucket Spreading Parallel Hash: A New, Robust, Parallel Hash Join Method for Data Skew in the Super Database Computer (SDC) 210
M Kitsuregawa and Y Ogawa

Session 7: Recovery and Data Replication

Two Epoch Algorithms for Disaster Recovery 222
H Garcia-Molina, R Hagmann and C Polyzois

Adaptable Recovery Using Dynamic Quorum Assignments 231
B Bhargava and S Browne

The Tree Quorum Protocol: An Efficient Approach for Managing Replicated Data 243
D Agrawal and A El Abbadi

Session 8: Deductive Databases 1

Factoring Augmented Regular Chain Programs 255
P T Wood

The Magic of Duplicates and Aggregates 264
I S Mumick, H Pirahesh and R Ramakrishnan

How to Forget the Past Without Repeating It 278
J F Naughton and R Ramakrishnan

Session 9: Query Optimization

Advanced Query Processing in Object Bases Using Access Support Relations . 290
A Kemper and G Moerkotte

**Elimination of Views and Redundant Variables in an SQL-like Database Language
for Extended NF^2 Structures** 302
N Südkamp and V Linnemann

Measuring the Complexity of Join Enumeration in Query Optimization 314
K Ono and G M Lohman

Session 10: Deductive Databases 2

Hybrid Transitive Closure Algorithms 326
R Agrawal and H V Jagadish

Distributed Transitive Closure Computations: The Disconnection Set Approach . 335
M A W Houtsma, P M G Apers and S Ceri

A Parallel Strategy for Transitive Closure Using Double Hash-Based Clustering . .347
J P Cheiney and C de Maindreville

Session 11: Deductive Databases 3

Rule Ordering Bottom-Up Fixpoint Evaluation of Logic Programs 359
R Ramakrishnan, D Srivastava and S Sudarshan

Efficient Implementation of Loops in Bottom-Up Evaluation of Logic Queries . . .372
J Kuittinen, O Nurmi, S Sippu and E Soisalon-Soininen

Right-, Left-, and Multi-Linear Rule Transformations that Maintain Context Information380
D Kemp, K Ramamohanarao and Z Somogyi

Session 12: Transaction Processing 2

ARIES/KVL: A Key-Value Locking Method for Concurrency Control of Multiaction Transactions Operating on B-Tree Indexes 392
C Mohan

Commit-LSN: A Novel and Simple Method for Reducing Locking and Latching in Transaction Processing Systems 406
C Mohan

The Effect of Skewed Data Access on Buffer Hits and Data Contention in a Data Sharing Environment 419
A Dan, D M Dias and P S Yu

Session 13: Conceptual Issues in OODBs

Consistency of Versions in Object-Oriented Databases	432
W Cellary and G Jomier	
Database Application Development as an Object Modeling Activity	442
M Jeusfeld, M Mertikas, I Wetzels, M Jarke and J W Schmidt	
ILOG: Declarative Creation and Manipulation of Object Identifiers	455
R Hull and M Yoshikawa	

Session 14: Database Machines

Tradeoffs in Processing Complex Join Queries via Hashing in Multiprocessor Database Machines	469
D A Schneider and D J DeWitt	
Hybrid-Range Partitioning Strategy: A New Declustering Strategy for Multiprocessor Database Machines	481
S Ghandeharizadeh and D J DeWitt	
An Adaptive Data Placement Scheme for Parallel Database Computer Systems . .	493
K A Hua and C Lee	

Session 15: Multidatabases and OODBs

A Multidatabase Transaction Model for InterBase	507
A K Elmagarmid, Y Leu, W Litwin and M Rusinkiewicz	
A Polygon Model for Heterogeneous Database Systems: The Source Tagging Perspective	519
Y R Wang and S E Madnick	
The C-based Database Programming Language Jasmine/C	539
M Aoshima, Y Izumida, A Makinouchi, F Suzuki and Y Yamane	

Session 16: Updates and Constraints

Synthesizing Database Transactions	552
X Qian	
Deriving Production Rules for Constraint Maintenance	566
S Ceri and J Widom	
Referential Integrity Revisited: An Object-Oriented Perspective	578
V M Markowitz	

Session 17: Spatial Databases

The Buddy Tree: An Efficient and Robust Access Method for Spatial Data Base Systems	590
B Seeger and H P Kriegel	
Query Processing for Distance Metrics	602
T L Wang and D Shasha	
On Indexing Line Segments	614
H V Jagadish	

Sessions 18: Nested Relations and Knowledge Based Systems

On Restructuring Nested Relations in Partitioned Normal Form	626
G Hulin	
Non-Monotonic Knowledge Evolution in VLKDBs	638
C Esculier	
Database Updates through Abduction	650
A C Kakas and P Mancarella	

Session 19: System Implementation Issues

An Incremental Join Attachment for Starburst 662
M Carey, E Shekita, G Lapis, B Lindsay and J McPherson

Distributed Linear Hashing and Parallel Projection in Main Memory Databases . . 674
C Severance, S Pramanik and P Wolberg

Efficient Main Memory Data Management Using the DBGraph Storage Model . . . 683
P Pucheral, J M Thevenin and P Valduriez

Session 20: Database Languages and Algorithms

A Probabilistic Framework for Vague Queries and Imprecise Information in Databases 696
N Fuhr

Priority-Hints: An Algorithm for Priority-Based Buffer Management 708
R Jauhari, M J Carey and M Livny

Concept Description Language for Statistical Data Modeling 722
T Catarci, G D'Angiolini and M Lenzerini