

CONTENTS

Session 1: Key-note speech

- From a laguna to open waters: Another view on the next generations of databases* 1
H. Gallaire

Session 2A: Panel Session

- Knowledge to Mediate from User's Workstations to Databases* 2

Session 2B: Data Models and Modeling

- On the design and implementation of information systems from deductive conceptual models* 3
A. Olive
- A deductive method for entity-relationship modeling* 13
G. Di Battista and M. Lenzerini
- A family of incomplete relational database models* 23
A. Ola and G. Ozsoyoglu

Session 3A: Extensible Databases and Data Structures

- Gral: an extensible relational database system for geometric applications* 33
R.H. Güting
- The LSD tree: spatial access to multidimensional point and non-point objects* 45
A. Henrich, H-W. Six and P. Widmayer
- Managing Complex objects in an extensible relational DBMS* 55
G. Gardarin, J-P. Cheiney, G. Kiernan, D. Pastre, and H. Stora

Session 3B: Parallelism

- Effective resource utilization for multiprocessor join execution* 67
M.C. Murphy and D. Rotem
- Optimization and dataflow algorithms for nested tree queries* 77
M. Muralikrishna
- Parallel processing of recursive queries in distributed architectures* 87
G. Hulin

Session 4A: Graphical Interfaces

- Pasta-3's graphical query language: direct manipulation, cooperative queries, full expressive power* 97
M. Kuntz and R. Melchert
- ENIAM: a more complete conceptual schema language* 107
P.N. Creasy

<i>FaceKit: a database interface design toolkit</i>	115
R. King and M. Novak	
 Session 4B: Parallelism	
<i>A low communication sort algorithm for a parallel database machine</i>	125
R.A. Lorie and H.C. Young	
<i>Percentile finding algorithm for multiple sorted runs</i>	135
B.R. Iyer, G.R. Ricard, and P.J. Varman	
<i>A signature access method for the Starburst Database System</i>	145
W.W. Chang and H.J. Schek	
 Session 5A: Recursive query optimization	
<i>Commutativity and its role in the processing of linear recursion</i>	155
Y.E. Ioannidis	
<i>Estimating the size of generalized transitive closures</i>	165
R.J. Lipton and J.F. Naughton	
<i>Argument Reduction by Factoring</i>	173
J.F. Naughton, R. Ramakrishnan, Y. Sagiv, and J.D. Ullman	
 Session 5B: Panel Session	
<i>Database support for hypertext</i>	183
 Session 6A: Recursive query optimization	
<i>Finding regular simple paths in graph databases</i>	185
A.O. Mendelzon and P.T. Wood	
<i>Towards an open architecture for LDL</i>	195
D. Chimenti, R. Gamboa and R. Krishnamurthy	
 Session 6B: Temporal Databases	
<i>Event-join optimization in temporal relational databases</i>	205
A. Segev and H. Gunadhi	
<i>Achieving zero information-loss in a classical database environment</i>	217
G. Bhargava and S.K. Gadia	
 Session 7A: Derived data and constraints	
<i>Derived data update in semantic databases</i>	225
I.A. Chen and D. McLeod	

<i>Using integrity constraints to provide intensional answers to relational queries</i>	237
A. Motro	
 Session 7B: Allocation and Optimization	
<i>Integration of buffer management and query optimization in relational database environment</i>	247
D.W. Cornell and Ph.S. Yu	
<i>The effect of bucket size tuning in the dynamic hybrid GRACE hash join method</i>	257
M. Kitsuregawa, M. Nakayama, and M. Takagi	
 Session 8A: Panel Session	
<i>Building knowledge-based applications with cooperating databases</i>	267
 Session 8B: Statistics and Statistical Databases	
<i>Random sampling from B^+ trees</i>	269
F. Olken and D. Rotem	
<i>Aggregate evaluability in statistical databases</i>	279
F.M. Malvestuto and M. Moscarini	
<i>Aggregates in Possibilistic Databases</i>	287
E.A. Rundensteiner and L. Bic	
 Session 9A: Complex Objects	
<i>Extending the relational algebra to capture complex objects</i>	297
B. Mitschang	
<i>Sorting, grouping and duplicate elimination in the advanced information management prototype</i>	307
G. Saake, V. Linnemann, P. Pistor, and L. Wegner	
<i>Optimization of relational schemas containing inclusion dependencies</i>	317
M.A. Casanova, L. Tucheran, A.L. Furtado, and A.P. Braga	
 Session 9B: Recovery and Concurrency Control	
<i>The case for Safe RAM</i>	327
G. Copeland, T. Keller, R. Krishnamurthy, and M. Smith	
<i>ARIES/NT: a recovery method based on write-ahead logging for nested transactions</i>	337
K. Rothermel and C. Mohan	
<i>Quasi Serializability: a correctness criterion for global concurrency control in InterBase</i>	347
W. Du and A. Elmagarmid	
 Session 10A: Object Management	
<i>The O_2 object manager: an overview</i>	357
F. Velez, G. Bernard, and V. Darnis	

<i>On correctly configuring versioned objects</i>	367
R. Agrawal and H.V. Jagadish	
<i>The Starburst Long Field Manager</i>	375
T.J. Lehman and B.G. Lindsay	
Session 10B: Priority Scheduling	
<i>Scheduling real-time transactions with disk resident data</i>	385
R. Abbott and H. Garcia-Molina	
<i>Priority in DBMS resource scheduling</i>	397
M.J. Carey, R. Jauhari, M. Livny	
Session 11A: Languages for OODB	
<i>The O₂ database programming language</i>	411
C. Lécluse and P. Richard	
<i>A model of queries for object-oriented databases</i>	423
W. Kim	
<i>OQL: a query language for manipulating object-oriented databases</i>	433
A.M. Alashqur, S.Y.W. Su, and H. Lam	
Session 11B: Panel Session	
<i>Database Tools and Interfaces</i>	443
Session 12A: Active databases	
<i>Monitoring database objects</i>	445
T. Risch	
<i>Situation monitoring for active databases</i>	455
A. Rosenthal, U.S. Chakravarthy, B. Blaustein, and J. Blakely	
Session 12B: Panel Session	
<i>Future research directions: Evidence from this conference</i>	465