

Proceedings

Tenth International Conference on

DATA ENGINEERING

February 14-18, 1994
Houston, Texas

Sponsored by
IEEE Computer Society Technical Committee on Data Engineering



IEEE Computer Society Press



The Institute of Electrical and Electronics Engineers, Inc.

Proceedings

**THE 10TH INTERNATIONAL
CONFERENCE ON
DATA ENGINEERING**

Dedication

The conference organizers would like to dedicate these proceedings to C.V. Ramamoorthy in recognition of his efforts in guiding the International Conference on Data Engineering continuously for the last 10 years.

Proceedings

THE 10TH INTERNATIONAL CONFERENCE ON DATA ENGINEERING

February 14 – 18, 1994

Houston, Texas

Sponsored by
IEEE Computer Society Technical Committee on Data Engineering



IEEE Computer Society Press
Los Alamitos, California

Washington • Brussels • Tokyo

The papers in this book comprise the proceedings of the meeting mentioned on the cover and title page. They reflect the authors' opinions and, in the interests of timely dissemination, are published as presented and without change. Their inclusion in this publication does not necessarily constitute endorsement by the editors, the IEEE Computer Society Press, or the Institute of Electrical and Electronics Engineers, Inc.



Published by the
IEEE Computer Society Press
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1264

© 1994 by the Institute of Electrical and Electronics Engineers, Inc. All rights reserved.

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limits of US copyright law, for private use of patrons, those articles in this volume that carry a code at the bottom of the first page, provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923. For other copying, reprint, or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331.

IEEE Computer Society Press Order Number 5400-02
IEEE Catalog Number 94CH3383-7
ISBN 0-8186-5400-7 (paper)
ISBN 0-8186-5401-5 (microfiche)
ISBN 0-8186-5402-3 (case)
ISSN 1063-6382

Additional copies can be ordered from

IEEE Computer Society Press
Customer Service Center
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1264
Tel: (714) 821-8380
Fax: (714) 821-4641
Email: cs.books@computer.org

IEEE Service Center
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331
Tel: (908) 981-1393
Fax: (908) 981-9667

IEEE Computer Society
13, avenue de l'Aquilon
B-1200 Brussels
BELGIUM
Tel: +32-2-770-2198
Fax: +32-2-770-8505

IEEE Computer Society
Ooshima Building
2-19-1 Minami-Aoyama
Minato-ku, Tokyo 107
JAPAN
Tel: +81-3-3408-3118
Fax: +81-3-3408-3553

Production Editor: Robert Werner
Cover art production by Joseph Daigle/Schenk-Daigle Studios
Printed in the United States of America by Braun-Brumfield, Inc.



The Institute of Electrical and Electronics Engineers, Inc.

Contents

Dedication	ii
Message from the General Chairs	x
Message from the Program Chair	xi
Organizing Committee	xii
List of Additional Reviewers	xiv
Steering Committee	xvi
Keynote Speech: Engineering Universal Access to Distributed Interactive Multimedia Data — Al Aho	
 Session 1: Analytical Modeling	
On a More Realistic Lock Contention Model and Its Analysis	2
A. Thomasian	
Approximate Analysis of Real-Time Database Systems	10
J.R. Haritsa	
Object Allocation In Distributed Databases and Mobile Computers	20
Y. Huang and O. Wolfson	
 Session 2: Data Replication and Distribution	
Supporting Partial Data Accesses to Replicated Data	32
P. Triantafillou and F. Xiao	
Efficient Support for Partial Write Operations in Replicated Databases	43
M. Rabinovich and E.D. Lazowska	
Mariposa: A New Architecture for Distributed Data	54
M. Stonebraker, P.M. Aoki, R. Devine, W. Litwin, and M. Olson	
 Session 3: Query Optimization	
Exploiting Uniqueness in Query Optimization	68
G.N. Paulley and P.-Å. Larson	
A Multi-Set Extended Relational Algebra: A Formal Approach to a Practical Issue	80
P.W.P.J. Grefen and R.A. de By	
Performing Group-By Before Join	89
W.P. Yan and P.-Å. Larson	
 Session 4: Disk Storage Management	
Analysis of Reorganization Overhead in Log-Structured File Systems	102
J.T. Robinson and P.A. Franaszek	
Performance Analysis of RAID5 Disk Arrays with a Vacationing Server Model for Rebuild Mode Operation	111
A. Thomasian and J. Menon	
Storage Reclamation and Reorganization in Client-Server Persistent Object Stores	120
V.-F. Yong, J.F. Naughton, and J.-B. Yu	

Storage Reclamation and Reorganization in Client-Server Persistent Object Stores V.-F. Yong, J.F. Naughton, and J.-B. Yu	120
 Session 5: Multidatabase Systems	
Cooperative Problem Solving Using Database Conversations T. Kirsche, R. Lenz, T. Ruf, and H. Wedekind	134
A Query Sampling Method for Estimating Local Cost Parameters in a Multidatabase System Q. Zhu and P.-Å. Larson	144
Resolving Attribute Incompatibility in Database Integration: An Evidential Reasoning Approach E.-P. Lim, J. Srivastava, and S. Shekhar	154
 Session 6: Panel Discussion	
The Role of Knowledge Mining in the Development and Evolution of New Applications Moderator: David Cohen Panel members: L. Berke, P. Bloom, D. Cohen, and D. Tsur	166
 Keynote Speech: From Data Engineering to Information Engineering Gio Wiederhold	
 Session 7: Panel Discussion	
The Impact of Database Research on Industrial Products Moderator: Daniel Barbará Panel members: J. Blakeley, D. Fishman, D. Lomet, C. Mohan, and M. Stonebraker	168
 Session 8: Schema Evolution	
Managing Change in the Rufus System P. Schwarz and K. Shoens	170
Polymorphic Reuse Mechanisms for Object-Oriented Database Specifications L. Liu, R. Zicari, K. Lieberherr, and W. Hürsch	180
A Method for Transforming Relational Schemas Into Conceptual Schemas P. Johannesson	190
 Session 9: Knowledge Bases and Rule Processing	
On the Interaction Between ISA and Cardinality Constraints D. Calvanese and M. Lenzerini	204
Comparing and Synthesizing Integrity Checking Methods for Deductive Databases C. García, M. Celma, L. Mota, and H. Decker	214
Discovering Database Summaries Through Refinements of Fuzzy Hypotheses D.H. Lee and M.H. Kim	223

Session 10: Parallel Databases

Declustering Techniques for Parallelizing Temporal Access Structures	232
V. Kouramajian, R. Elmasri, and A. Chaudhry	
Disk Allocation Methods for Parallelizing Grid Files	243
Y. Zhou, S. Shekhar, and M. Coyle	
Object Placement in Parallel Object-Oriented Database Systems	253
S. Ghandeharizadeh, D. Wilhite, K. Lin, and X. Zhao	

Session 11: Temporal Databases

Implementing Calendars and Temporal Rules in Next Generation Databases	264
R. Chandra, A. Segev, and M. Stonebraker	
The TP-Index: A Dynamic and Efficient Indexing Mechanism for Temporal Databases	274
H. Shen, B.C. Ooi, and H.J. Lu	
Efficient Evaluation of the Valid-Time Natural Join	282
M.D. Soo, R.T. Snodgrass, and C.S. Jensen	

Session 12: Knowledge Mining

Supporting High-Bandwidth Navigation in Object Bases	294
V. Vasudevan	
Supporting Data Mining of Large Databases by Visual Feedback Queries	302
D.A. Keim, H.-P. Kriegel, and T. Seidl	
QBISM: Extending a DBMS to Support 3D Medical Images	314
M. Arya, W. Cody, C. Faloutsos, J. Richardson, and A. Toga	

Session 13: Indexing Techniques

Efficient Organization of Large Multidimensional Arrays	328
S. Sarawagi and M. Stonebraker	
Index Structures for Information Filtering Under the Vector Space Model	337
T.W. Yan and H. Garcia-Molina	
An Efficient Relational Implementation of Recursive Relationships Using Path Signatures	348
J. Teuhola	
Performance Evaluation of Grid Based Multi-Attribute Record Declustering Methods	356
B. Himatsingka and J. Srivastava	

Session 14: Technology and Applications

Knowledge-Based Handling of Design Expertise	368
P. Morizet-Mahoudeaux, E. Suzuki, and S. Ohsuga	
Active Databases for Active Repositories	375
H. Jasper	
Papyrus: A History-Based VLSI Design Process Management System	385
T. Chiueh and R. Katz	
X.500 Directory Schema Management	393
J.W. Hong, D.L. Silver, and M.A. Bauer	

Session 15: Invited Industrial Presentations I

Data Management in Delayed Conferencing	402
A. Hsu	
Transactional Workflows: Research, Enabling Technologies, and Applications	403
A. Sheth	
Parallel Approaches to Database Management	404
D. Reiner	

Session 16: Query Processing

Sort-Merge-Join: An Idea Whose Time Has(h) Passed?	406
G. Graefe	
Data Placement and Buffer Management for Concurrent Mergesorts with Parallel Prefetching	418
K.-L. Wu, P.S. Yu, and J.Z. Teng	
Fast Ranking in Limited Space	428
A. Moffat and J. Zobel	

Session 17: Heterogeneity and Multidatabase Systems

Capturing Design Dynamics — The CONCORD Approach	440
N. Ritter, B. Mitschang, T. Härder, M. Gesmann, and H. Schöning	
Semantics-Based Multilevel Transaction Management in Federated Systems	452
A. Deacon, H.-J. Schek, and G. Weikum	
Specification and Management of Extended Transactions in a Programmable Transaction Environment	462
D. Georgakopoulos, M. Hornick, P. Krychniak, and F. Manola	

Session 18: Panel Discussion

Nomadic, Wireless Computing

Moderator: T. Imielinski

Panel members: R. Alonso, W. Litwin, M. Stonebraker, and Y. Wolfson

Session 19: Advanced Query Processing Techniques

Query Optimization Strategies for Browsing Sessions	478
M.L. Kersten and M.F.N. de Boer	
Analysis of Common Subexpression Exploitation Models in Multiple-Query Processing	488
J.R. Alsabbagh and V.V. Raghavan	
A Hybrid Transitive Closure Algorithm for Sequential and Parallel Processing	498
Q. Yang, C. Yu, C. Liu, T. Pham, S. Dao, and G. Wang	

Session 20: Object-Oriented Databases

<i>Object Skeletons: An Efficient Navigation Structure for Object-Oriented Database Systems</i>	508
K. Hua and C. Tripathy	
Applying Signatures for Forward Traversal Query Processing in Object-Oriented Databases	518
H.S. Yong, S. Lee, and H.-J. Kim	
On the Selection of Optimal Index Configuration in Object-Oriented Databases	526
S. Choenni, E. Bertino, H.M. Blanken, and T. Chang	

Session 21: Invited Industrial Presentations II

Transactional Workflow Management in Distributed Object Computing Environments	540
D. Georgakopoulos	
Distributed Heterogeneous Information Systems	541
W.A. Carpenter, G. Emami, E.G. Mamas, and A.C. Nigam	

Session 22: Plenary Panel Discussion

Moderator: M. Stonebraker
Panelists: To be announced

Index of Authors	543
-------------------------	------------

Message from the General Chairs

Welcome to the Tenth International Conference on Data Engineering. This year we have arrived at a milestone that attests to the vision of the founders of this special meeting. To start planning over ten years ago a major conference devoted to an area overlapping the interests of academicians and practitioners is a feat that could only be accomplished by visionaries such as those listed on our steering committee. We pay special tribute to C.V. Ramamoorthy for guiding this endeavor continuously through the last 10 years.

As usual the conference will last for five days. The first two are devoted to tutorials, the last three are devoted to technical presentations, panels and industrial sessions. This year the conference also features a full day tutorial on the Sunday before the official start of the conference. This tutorial is held in conjunction with the RIDE workshop series.

We are fortunate to have among us two distinguished computer scientists. Alfred Aho will give the keynote address on the opening day of the conference. Gio Wiederhold will give the keynote address of the second day. Gio's relationship to ICDE goes back to the inception of this conference. He was the program chair of the second ICDE held in Los Angeles, California.

The technical program of the conference was selected by a distinguished program committee. The committee, led by 10 vice chairs, all worked together to insure a first rate program. The credit for selecting, maintaining and conducting the program committee business was all left in the able hands of our esteemed friend and colleague, Marek Rusinkiewicz. Marek has diligently worked on the selection of these technical papers. A million thanks for a job well done.

Special thanks to Omran Bukhres and Peter Muth for a first rate tutorial program. Not unexpectedly, Witold Litwin was successful in selecting three excellent panels. Mike Stonebraker is responsible for selecting the best paper. Local arrangements were handled by J.F. Paris, C. Eick and A. Cheng. European coordination was done by our good friends Eva Kuehn and Gunter Schlageter. Far east coordination was done by Yutaka Matsushita and Mike Papazoglou. Finally, I would like to single out with special thanks Ernst Leiss, our treasurer, and Sumi Helal, for taking care of all the publicity. Sumi has shown tremendous dedication to the conference.

We hope that you join our celebration of the tenth anniversary of the conference in the usual Texan hospitality. The local arrangements folks have done a great job in selecting social events that will get even those with two left feet on the dance floor. We are planning a little birthday party for ICDE during the banquet which will be hosted by the steering committee. Please come join us in Houston for a week of intellectual pursuit, fun and renewals of friendships.

Erich Neuhold and Ahmed Elmagarmid
General Co-Chairs

Message from the Program Chair

These proceedings contain papers presented at the Tenth International Conference on Data Engineering. This year we received 270 submissions from 31 countries. After a thorough review the program committee decided to accept 50 papers. In addition to research paper sessions, the technical program of the conference includes invited industrial presentations, panel discussions, and keynote addresses by Al Aho and Gio Wiederhold.

It is a pleasant duty to express my sincere thanks to all who helped to create an interesting technical program, and especially:

- Authors of all papers submitted to ICDE-10;
- vice-chairs of the program committee, Elisa Bertino, Ramez Elmasri, Tomasz Imielinski, Masaru Kitsuregawa, Paul Larson, Dennis McLeod, C. Mohan, Hans Schek, and Richard Snodgrass for their time and effort in preparing the program of the conference;
- all members of the program committee, as well as 202 additional reviewers of the papers;
- Amit Sheth, who coordinated the Technology and Application Track, Daniel Barbara, who put together the Industrial Program, and Witold Litwin, who was responsible for the panel discussions;
- Ahmed Elmagarmid and Erich Neuhold, the general co-chairs of the conference who allowed us to concentrate on the technical program by taking care of all remaining problems; and
- members of the steering committee for their laconic but sage advice.

My special thanks go to Nancy Tran, George Karabatis, Andrzej Cichocki and Paolo Missier at the University of Houston, without whom it would have been impossible to complete the program committee work in time.

We hope that the conference will be stimulating and enjoyable to all who attend it, and keeping with the tradition of previous years, the papers in the proceedings will be both interesting and relevant.

Welcome y'all to Texas!

Marek Rusinkiewicz
Program Chair

Organizing Committee

General Chairs

Ahmed K. Elmagarmid, Purdue Univ. and Erich Neuhold, GMD-IPSI Darmstadt

Program Chair

Marek Rusinkiewicz, Univ. of Houston

Tutorial Program

Omran Bukhres, Purdue Univ. and Peter Muth, GMD-IPSI Darmstadt

Panel Program

Witold Litwin, Univ. of California at Berkeley

Program Committee

Rafael Alonso, MITL

B. R. Badrinath, Rutgers Univ.

Elisa Bertino, Univ. of Genoa, Vice Chair — Object-Oriented Systems

Yuri Breitbart, Univ. of Kentucky

Alex Buchmann, TH Darmstadt

Phil Canatta, MCC

John Carlis, Univ. of Minnesota, Publications

Arbee Chen, National Tsing Hua Univ.

Peter Dadam, Univ. Ulm

Amr El Abbadi, Univ. California at Santa Barbara

Ramez Elmasri, Univ Texas at Arlington, Vice Chair — Data Modeling

George Gardarin, INRIA

Dimitrios Georgakopoulos, GTE Labs

Goetz Graefe, Portland State Univ.

Ralf Hartmut Gueting, Univ. Hagen

Theo Haerder, Univ. Kaiserslautern

Tomasz Imielinski, Rutgers Univ., Vice Chair — Database Theory

Bala Iyer, IBM DB Technology Lab, San Jose

H.V. Jagadish, AT&T Bell Labs

Matthias Jarke, Aachen Univ.

Christian Jensen, Aalborg University

Jie-Yong Juang, National Taiwan Univ.

Leonid Kalinichenko, Russian Academy of Sciences

Yahiko Kambayashi, Kyoto Univ.

Larry Kerschberg, George Mason Univ.

Masaru Kitsuregawa, Vice Chair — Parallel Databases

Yasushi Kiyoki, Univ. of Tsukuba

Wolfgang Klas, GMD-IPSI Darmstadt

Eva Kuehn, University of Technology Vienna

Paul Larson, Waterloo Univ., Vice Chair — Access Methods

Dik Lee, Ohio State Univ.

Witold Litwin, Univ. of California at Berkeley

Akifumi Makinouchi, Kyushu Univ.

Yutaka Matsushita, Keio Univ.
James McKenna, Bellcore
Dennis McLeod, UCS, Vice Chair — Heterogeneous Databases
Alberto Mendelzon, Univ. Toronto
Michele Missikoff, IASI-CNR, Rome
C. Mohan, IBM Almaden, Vice Chair — Transaction Management
John Mylopoulos, Univ. of Toronto
Mike Papazoglou, Queensland Univ. of Technology
J.F. Paris, University of Houston
Calton Pu, Columbia Univ.
Raghu Ramakrishnan, Univ. of Wisconsin
Krithi Ramamritham, Univ. of Massachusetts
Andreas Reuter, Univ. Stuttgart
Nick Roussopoulos, Univ. of Maryland
Hans Schek, ETH Zurich, Vice Chair — Extensible and Spatial Databases
Peter Scheuermann, Northwestern Univ.
Gunter Schlageter, Univi. of Hagen
Arie Segev, Univ of California at Berkely
Ming Shan, HP Labs
Amit Sheth, Technology and Applications track
Anoop Singhal, AT&T Labs
Richard Snodgrass, Univ. of Arizona, Vice Chair — Temporal Databases
Mike Stonebraker, Univ. California at Berkeley (Awards)
Alex Tomassian, IBM Watson
Susan Urban, Arizona State Univ.
Gerhard Weikum, ETH Zurich
Antoni Wolski, Research Center of Finland
Gene Wu, Bellcore

European Coordinators

Eva Kuehn, University of Technology Vienna
Gunter Schlageter, University of Hagen

Far East Coordinators

Yutaka Matsushita, Keio University
Mike Papazoglou, Queensland University of Technology

Industrial and Exhibits Program

Daniel Barbara, MITL

Publicity

Abdelsalam Helal, University of Texas at Arlington

Financial Chair

Ernst Leiss, University of Houston

Local Arrangements

J.F. Paris, University of Houston, Chair
C. Eick, University of Houston
A. Cheng, University of Houston, Registration

List of Additional Reviewers

Aberer, K.
Alonso, G.
Amano, H.
Aref, W.G.
Arikawa, M.
Attaluri, G.
Baekgaard, L.
Becker, W.
Berkel, T.
Bestavros, A.
Beuter, T.
Biliris, A.
Boecker, D.
Bouguettaya, A.
Bradshaw, D.
Breiteneder, C.
Brinkhoff, T.
Burkowski, F.
Busse, R.
Cambray de, B.
Castano, S.
Chang, S.K.
Chen, Q.
Choenni, S.
Chong, E.I.
Christodoulakis, S.
Chu, W.
Coburn, N.
Davison, D.L.
Dellis, A.
Dessloch, S.
Devine, B.
Du, W.
Dyballa, A.
Dyreson, C.
Eitler, M.
Farrell, R.
Feldkamp, J.
Fisher, G.
Franklin, M.
Fujita, K.
Furtado, A. L.
Furukawa, T.
Gassner, K.
Gertz, M.
Geuder, U.
Grebe, J.
Guenther, O.
Guenthoer, R.
Guerrini, G.
Harada, L.
Hasse, C.
Heinlein, C.
Hirano, Y.
Hofri, M.
Hornick, M.
Houtsma, M.A.W.
Huhns, M.
Ichimura, S.
Inoue, U.
Iwaihara, M.
Jensen, F.
Kaefer, W.
Kakeshita, T.
Kalus, C.
Karabatis, G.
Kawakita, H.
Kessler, U.
Kiernan, G.
Kilov, H.
Kimura, Y.
Kogalovsky, M.
Kolovson, C.
Konomi, S.
Korth, H.
Kouramajian, V.
Krishnakumar, N.
Ksiezzyk, T.
Kutschera, P.
Kuznetsov, S.
Lavender, G.
Lee, C.
Lee, J.
Leichner, S.
Lenzen, F.
Leong, H.V.
Lien, Y.N.
Lieumen, D.
Lindenau, B.
Lipeck, U.
Liu, M.L.
Liu, X.
Ludwig, T.
Marcus, H.
Mark, L.
Masunaga, Y.
Matheus, C.
Meadows, C.
Mehrotra, S.
Meier, J.
Milo, T.
Miyazaki, N.
Missier, P.
Mitrach, S.
Moenkeberg, A.
Montesi, D.
Mueck, T.
Mullen, J.
Murphy, M.
Murthy, V.
Nishiyama, H.
Nakamura, M.
Nakano, M.
Norrie, M.
Novikov, B.
Obermarck, R.
Ogawa, Y.
Ohboshi, N.
Ohmori, T.
Ohsawa, Y.
Olken, F.
Omiecinski, E.
Ouksel, A.
Palomaki, A.
Paulley, G.
Perrizo, W.K.
Pitoura, E.
Plexousakis, D.
Pohlai, H.
Pollak, R.
Prodingler, H.
Rahm, E.
Reichert, M.
Resende, R.F.
Richter, J.P.
Ross, K.A.
Rotem, D.
Sakakibara, H.
Salzer, G.
Samarati, P.
Satoh, T.
Scherer, A.
Schmidt, U.
Schneider, D.A.

Schneider, R.
Scholtz, G.
Schwarz, K.
Schwenkreis, F.
Selfridge, P. G.
Semeczko, G.
Sha, G.
Shen, W.M.
Shimajo, S.
Siegel, M.
Silverstein, E.
Singh, M.
Smedt de, P.
Soo, M.
Srivastava, D.
Straube, M.
Sudarshan, S.
Sullivan, M.
Sun, W.
Suzuki, S.

Suzuki, T.
Swami, S.
Takakura, H.
Tari, Z.
Tomlinson, C.
Topaloglou, T.
Triantafillou, P.
Tschernko, T.
Tsangaris, M.M.
Tsotras, V.
Tsuchida, M.
Tsuruoka, K.
Tuzhilin, A.
Udagawa, Y.
Veijalainen, J.
Vingralek, R.
Waechter, H.
Wei, S.S.
Widmayer, P.
Widom, J.

Wiener, J.
Wilkes, C.T.
Winckler, A.
Woerner, B.
Wong, M.H.
Wood, P.
Wuu, S.Y.
Xu, H.
Yan, W.P.
Yeh, T.S.
Yokota, H.
Yoshikawa, A.
Yu, C.
Zabback, P.
Zadorozhny, V.
Zhu, Q.
Zink, R.
Ziu, X.

Steering Committee

P. Bruce Berra — Syracuse University

John Carlis — University of Minnesota

Nick Cercone — University of Regina

Tadao Ichikawa — Hiroshima University

Ming (Mike) T. Liu — Ohio State University

C.V. Ramamoorthy — University of California

Joseph Urban — Arizona State University

Benjamin Wah — University of Illinois