

The
VLDB Journal

The International Journal on Very Large Data Bases

Volume 2(4) (1993)

The **VLDB** Journal

The International Journal on Very Large Data Bases

Editors-in-Chief

Fred J. Maryanski
Storrs, CT, USA

Hans-J. Schek
Zürich, Switzerland

Editorial Board

Serge Abiteboul
Peter Buneman
Hector Garcia-Molina
Nathan Goodman
Kuan-Tsae Huang
Masaru Kitsuregawa
Dennis McLeod
Antoni Olivé
F.J. Radermacher
Ron Sacks-Davis
Joachim W. Schmidt
Stanley Y.W. Su
Hartmut Wedekind

Michel Adiba
Walter A. Burkhard
Georges Gardarin
Georg Gottlo
Tadao Ichikawa
Tosiyasu L. Kunii
Robert A. Meersman
M. Tamer Özsu
K. Ramamohanarao
Peter Scheuerman
Kenneth Sevcik
Yannis Vassiliou
Stanley Zdonik

Antonio Albano
Steven A. Demurjian
Hans-Detlef Gerhardt
Peter Gray
Roger King
Michel Léonard
John Mylopoulos
Alain Pirotte
Andreas Reuter
Gunter Schlageter
Arne Solvberg
Kyu-Young Whang

Publication Board

Arie Shoshani
Thomas Wu

Fred J. Maryanski

Hans-J. Schek
Michael Rys



Volume 2(4) (1993)

**THE BOXWOOD PRESS
PACIFIC GROVE, CA, USA**

© 1993 BY THE VLDB ENDOWMENT

Copyright reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission of the copyright owner.

Submission of a paper to *The VLDB Journal* is understood to imply that it is not being considered for publication elsewhere and that the author's permission to publish his/her article(s) in this journal implies the exclusive authorization of the publisher to deal with all issues concerning the copyright therein.

Submission of multi-authored manuscripts to this journal implies the consent of *each* of the authors. The publisher will assume that the senior or corresponding author has specifically obtained the approval of all other co-authors to submit the manuscript to this journal.

No responsibility is assumed by the Publisher for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions or ideas contained in the material herein.

Special regulations for authors. Upon acceptance of an article by the journal, the author(s) will be asked to transfer copyright of the article to the publisher. This transfer will ensure the widest possible dissemination of information.

Published quarterly by:

THE BOXWOOD PRESS

183 Ocean View Blvd.

Pacific Grove, CA 93950, USA

Telephone: (408) 375-9110/Fax: (408) 375-0430

Note to the Subject Index

Users of an index have specific topics in mind and wish to find cues to subject matter discussed in the volume. These cues (words and phrases) should refer to well-defined structures, actions, or concepts of the subject matter—the subject-matter organization or taxonomy. Because the study of very large databases is a relatively new academic pursuit, the terminology is evolving and sometimes has inconsistent or undefined referents. The Editors of *The VLDB Journal* are attempting to develop a useful, consistent vocabulary for the articles published in the journal. Key words alone may not be sufficient to encompass the content of the articles. Authors can help this effort by highlighting words or short phrases in their manuscripts that should be indexed in addition to the key words. New terms should be defined carefully when they are introduced. Readers can help significantly by writing to the Editor, suggesting more precise meanings of terms used in this subject matter.

Subject Index for Volumes 1(1) through 2(4)

- Algorithm, priority scheduling, 2(2):117
- Architecture, extensible 1(2):241
- Atomicity, 2(4):407
 - global, 1(2):181
 - multidatabase system, 1(1):1
- Autonomy, 1(2):181
- B+ -tree, 2(4):361
- Benchmark, test database generation, 2(2):173
- Buffer management, 2(1):1
 - simulation of, 2(1):1
- Compensation 1(2):181, 2(4):407
- Complex object transaction management, 2(4):407
- Concurrency, priority scheduling, 2(2):117
- Concurrency control, 2(1):39, 2(2):215, 2(3):331
 - B+ -tree, 2(4):361
 - global, 2(3):331
 - multidatabase system, 2(3):331
 - performance of, 2(4):361
 - simulation of, 2(4):361
- Consistency, updates, 2(2):215
- Consistency constraint, test data, 2(2):173
- Cooperation, 1(1):41
 - of database system, 2(3):331
- Cooperative application, 1(1):41
- Cooperative transaction hierarchy, 1(1):41
- Correctness, transaction, 1(1):41
- Data abstraction, semantic relationship, 2(4):455
- Data contention, 2(2):117, 2(4):361
- Data model,
 - entity relationship, 2(4):455
 - relational, 2(4):455
 - semantic, 2(4):455
- Data representation,
 - of imprecise data, 2(4):490
 - of partial value, 2(4):490
- Data semantics, 1(1):81
- Data sharing, 1(1):127, 1(2):285
- Data skew, 2(3):303
- Database conversion, 1(1):127, 1(2):285
- Database design, 2(2):173, 2(4):455
 - tool, 2(4):455
- Database integration, 1(1):81
- Database system,
 - cooperation, 2(3):331
 - development toolkit, 1(2):241
 - distributed, 1(1):81
 - federated, 1(1):81, 1(1):127, 1(2):285, 2(2):215
 - heterogeneous, 1(1):81, 1(1):127, 1(2):285, 2(2):241
 - interoperable, 1(2):241
 - multi-, 1(1):1, 1(2):181, 2(2):153, 2(2):215, 2(3):331
 - parallel, 2(3):277, 2(3):303
 - real-time, 2(2):117
 - temporal, 2(1):75
- Deadlock detection, 1(1):41
 - nested transaction, 2(1):39
- Deadlock recovery, 1(1):1
- Deadlock, global, 1(2):181
- Design transaction, 1(1):41
- Differential file, 2(1):75
- Distributed database system, 1(1):81
- Distributed information system, 2(3):243

Document information retrieval, 2(3):243
 Entity-relationship data model, 2(4):455
 Federated database system, 1(1):81, 1(1):127,
 1(2):285, 2(2):215
 Federated system architecture, 1(1):127,
 1(2):285
 File organization, 2(3):243
 File, inverted, 2(3):243
 Full text information retrieval, 2(3):243
 Global transaction scheduling, 2(3):331
 Graph matching, 2(4):490
 Heterogeneous database system, 1(1):81,
 1(1):127, 1(2):285
 integration of, 1(2):241
 Heterogeneous system, 1(1):127, 1(2):285
 Imprecise data, representation of, 2(4):490
 Index, inverted, 2(3):243
 Information retrieval,
 document, 2(3):243
 full text, 2(3):243
 Information system,
 distributed, 2(3):243
 parallel, 2(3):243, 2(3):277, 2(3):303
 Integration,
 database, 1(1):81
 heterogeneous database system, 1(2):241
 schema, 1(1):81
 Inter-operation parallelism, 2(3):277,
 2(3):303
 Inter-transaction parallelism, 2(4):407
 Interoperability, 1(1):127, 1(2):285, 2(2):153
 Interoperable database system, interface,
 1(2):241
 Intra-operation parallelism, 2(3):277,
 2(3):303
 Intra-query parallelism, 2(3):303
 Intra-transaction parallelism, 2(1):39,
 2(4):407
 Inverted file, 2(3):243
 Inverted index, 2(3):243
 Join method, 2(1):1
 Join, multi-way, 2(3):303
 Load balancing, dynamic, 2(3):303
 Locking,
 hierarchical, 2(1):39
 lock inheritance, 2(1):39
 lock mode, 2(1):39
 lock upgrading/downgrading, 2(1):39
 nested transaction, 2(1):39
 object hierarchy, 2(1):39
 strategy, 2(4):361
 Logging,
 multi-level, 2(4):407
 transaction, 1(1):1
 Model translation, 1(1):127, 1(2):285
 Multi-level logging, 2(4):407
 Multi-level recovery, 2(4):407
 Multi-level transaction, 2(4):407
 performance of, 2(4):407
 Multi-way join, 2(3):303
 Multidatabase language, 2(2):153
 Multidatabase system, 1(1):1, 1(2):181,
 2(2):153, 2(2):215, 2(3):331
 concurrency control, 2(3):331
 global concurrency control, 2(3):331
 Multirelation, 2(2):153
 algebra, 2(2):153
 calculus, 2(2):153
 Nested transaction, 1(1):41
 deadlock detection in, 2(1):39
 locking, 2(1):39
 transaction tree, 2(1):39
 Non-serializability, 1(1):41
 Object caching, 2(1):75
 Object hierarchy, locking of, 2(1):39
 Open nested transaction, 2(4):407
 Operator tree,
 bushy, 2(3):303
 leftdeep, 2(3):303
 rightdeep, 2(3):303
 zigzag, 2(3):277
 Parallel database system, 2(3):277,
 2(3):303
 Parallel information system, 2(3):243,
 2(3):277, 2(3):303
 Parallel query processing, 2(3):243,
 2(3):277
 Parallelism,
 inter-operation, 2(3):277, 2(3):303
 inter-transaction, 2(4):407
 intra-operation, 2(3):277, 2(3):303
 intra-query, 2(3):303
 intra-transaction, 2(1):39, 2(4):407
 shared-memory, 2(3):277
 shared-nothing, 2(3):243, 2(3):277,
 2(3):303
 Partial value,
 efficient communication of, 2(4):490
 minimal semantically equivalent sub-
 sets of, 2(4):490

- representation of, properties of, 2(4):490
- Persistence, 2(4):407
- Priority concurrency algorithm, 2(2):117
- Priority scheduling, 2(2):117
 - simulation of, 2(2):117
- Query execution,
 - differential file, 2(1):75
 - object caching, 2(1):75
- Query languages, 2(2):153
- Query optimization, 2(1):1, 2(2):153, 2(3):243, 2(3):277, 2(3):303
 - in multidatabase systems, 2(2):153
- Query processing,
 - with imprecise data, 2(4):490
 - parallel, 2(3):243, 2(3):277
- Queueing model, 2(1):1
- Real-time database system, 2(2):117
- Recovery,
 - deadlock, 1(1):1
 - multi-level, 2(4):407
 - multidatabase system, 1(1):1
 - strong, 1(2):181
- Redo, 1(2):181
- Relational data model, 2(4):455
- Replicated data management, 2(2):215
- Replication control, 2(2):215
- Resolvable conflicts, 2(2):215
- Resource contention, 2(2):117
- Resource management, 2(1):1
- Retry, 1(2):181
- Return on consumption, 2(1):1
- Scheduling,
 - deadline, 2(2):117
 - in real-time database systems, 2(2):117
 - value-based, 2(2):117
- Schema integration, 1(1):81
- Schema translation, 1(1):127, 1(2):285
- Semantic assertion, 1(1):81
- Semantic data model, 2(4):455
- Semantic heterogeneity, 1(1):81
- Semantic relationships, data abstraction, 2(4):455
- Serializability, 1(2):181, 2(2):215, 2(3):331
 - chain-conflicting, 2(3):331
 - epsilon, 1(2):181
 - global, 1(2):181, 2(3):331
 - quasi-, 1(2):181, 2(2):215
 - sharing, 2(3):331
 - strong, 1(2):181
 - two-level, 1(2):181
- Simulated annealing, 2(1):1
- Striping, 2(3):243
- Temporal database system, 2(1):75
- Test data, 2(2):173
 - consistency constraint, 2(2):173
 - generating consistent, 2(2):173
 - search space, 2(2):173
 - validation, 2(2):173
- Test database, 2(2):173
 - generation of, benchmark, 2(2):173
- Toolkit, database system development, 1(2):241
- Transaction
 - correctness, 1(1):41
 - design, 1(1):41
 - hierarchy, cooperative, 1(1):41
 - multi-level, 2(4):407
 - nested, 1(1):41
 - open nested, 2(4):407
 - processing, 2(2):215, 2(4):407
 - scheduling, 2(2):117
 - time, relational model, 2(1):75
 - translation, 1(1):127, 1(2):285
 - tree, nested transaction, 2(1):39
- Transaction management,
 - complex object, 2(4):407
 - integration protocol for, 1(2):241
 - in multidatabase systems, 1(1):1, 1(2):181
- Transaction priority, 2(2):117
 - performance, 2(2):117
- Translation,
 - model, 1(1):127, 1(2):285
 - schema, 1(1):127, 1(2):285
 - transaction, 1(1):127, 1(2):285
- Tree, B+, 2(4):361
- Two-phase commit, 1(2):181
- Update consistency, 2(2):215
- Validation, test data, 2(2):173
- Value-based scheduling, 2(2):117
- Version management, 1(1):41

Author Index for Volumes 1(1) through 2(4)

- Borla-Salamet, P., 2:277
Breitbart, Y., 1:1
Bukhres, O., 2:215
- Carey, M.J., 2:117, 2:361
Chen, A.L.P., 2:490
Cornell, D.W., 2:1
- Drew, P., 1:241
Du, W., 2:215
Dupont, Y., 1:81
- Elmagarmid, A.K., 2:215, 2:331
- Garcia-Molina, H., 1:181, 2:243
Grant, J., 2:153
- Härder, T., 2:39
Haritsa, J.R., 2:117
Hasse, C., 2:407
Heimbigner, D., 1:241
Hsiao, D., 1:127, 1:285
Hua, K.A., 2:303
- Jensen, C.S., 2:75
- Kim, W., 2:215
King, R., 1:241
- Litwin, W., 2:153
Livny, M., 2:117
Lo, Y.-L., 2:303
Lockemann, P.C., 2:173
- Mark, L., 2:75
Moerkotte, G., 2:173
- Neufeld, A., 2:173
Nodine, M., 1:41
- Parent, C., 1:81
- Rothermel, K., 2:39
Roussopoulos, N., 2:75, 2:153
- Sellis, T., 2:75
Silberschatz, A., 1:1, 1:181
Spaccapietra, S., 1:81
Srinivasan, V., 2:361
Storey, V.C., 2:455
- Thompson, G.R., 1:1
Tomasic, A., 2:243
Tseng, F.S.C., 2:490
- Weikum, G., 2:407
- Yang, W.-P., 2:490
Young, H.C., 2:303
Yu, P.S., 2:1
- Zait, M., 2:277
Zdonik, S., 1:41
Zhang, A., 2:331
Ziane, M., 2:277

Reviewers for Volumes 1(1) through 2(4)

We wish to acknowledge the important assistance we have received from the following individuals who evaluated manuscripts submitted for publication from 1991 to 1993.

Abiteboul, S.
Adiba, M.
Albano, A.
Atzeni, P.

Batini, C.
Batory, D.S.
Beeri, C.
Berrut, C.
Bertino, E.
Bhargava, B.
Biskup, J.
Blanken, H.
Borgida, A.
Bouzeghoub, M.
Buneman, P.
Burkhard, W.

Chu, W.
Cohen, P.R.

Dadam, P.
Dayal, U.
Demurjian, S.A.
Deshpande, A.
DeWitt, D.
Dittrich, K.R.
Du, W.

Effelsberg, W.
Ehrich, H.-D.
Elmagarmid, A.K.

Faloutsos, C.

Garcia-Molina, H.
Gardarin, G.
Georgakopoulos, D.
Ghelli, G.
Gottlob, G.
Gray, M.D.

Härder, T.
Hagelstein, J.
Halatsis, C.
Houtsma, M.A.W.
Hsu, M.

Ichikawa, T.

Jagadish, H.V.

Kambayashi, Y.
King, R.
Kitsuregawa, M.
Kramer, R.
Küspert, K.
Kunii, T.L.

Larson, P.-A.
Lausen, G.
Lee, D.
Leonard, M.
Levy, E.
Liu, L.
Lochovsky, F.
Lockemann, P.
Lu, H.

Mannila, H.
Manthey, R.
McLeod, D.
Meersman, R.
Moens, M.
Moss, E.
Murphy, M.
Mylopoulos, J.

Navathe, S.B.

Özsu, M.T.
Olive, A.

Pirotte, A.

Pistor, P.
Pylyshyn, Z.

Radermacher, F.J.
Ramamohanarao, K.
Reuter, A.
Rothermel, K.

Sacks-Davis, R.
Scheurman, P.
Schlageter, G.
Schmidt, J.W.
Schneider, D.A.
Scholl, M.
Scholl, M.A.
Scholz, G.
Schwarz, P.
Selinger, P.
Sellis, T.
Sevcik, K.
Silberschatz, A.

Solvberg, A.
Spaccapietra, S.
Stemple, D.
Stocker, P.
Su, S.

Thompson, G.R.
Tiberio, P.

Valduriez, P.
Van Gucht, D.
Vassiliou, Y.

Wahlster, W.
Walter, B.
Wedekind, H.
Weikum, G.
Whang, K.-Y.
Wisdom, J.

Zdonik, S.