

Lecture Notes in Computer Science

826

Edited by G. Goos and J. Hartmanis

Advisory Board: W. Brauer D. Gries J. Stoer



David S. Bowers (Ed.)

Directions in Databases

12th British National Conference on Databases,
BNCOD 12
Guildford, United Kingdom, July 6-8, 1994
Proceedings

Springer-Verlag

Berlin Heidelberg New York
London Paris Tokyo
Hong Kong Barcelona
Budapest

Series Editors

Gerhard Goos
Universität Karlsruhe
Postfach 69 80
Vincenz-Priessnitz-Straße 1
D-76131 Karlsruhe, Germany

Juris Hartmanis
Cornell University
Department of Computer Science
4130 Upson Hall
Ithaca, NY 14853, USA

Volume Editor

David S. Bowers
Department of Mathematical and Computing Sciences, University of Surrey
Guildford GU2 5XH, United Kingdom

CR Subject Classification (1991): H.2, H.4

ISBN 3-540-58235-5 Springer-Verlag Berlin Heidelberg New York
ISBN 0-387-58235-5 Springer-Verlag New York Berlin Heidelberg

CIP data applied for

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer-Verlag. Violations are liable for prosecution under the German Copyright Law.

© Springer-Verlag Berlin Heidelberg 1994
Printed in Germany

Typesetting: Camera-ready by author
SPIN: 10472592 45/3140-543210 - Printed on acid-free paper

Preface

This volume continues the theme established for the series of British National Conferences on Databases, and presents a number of papers which address current *Directions in Databases*. The papers themselves constitute the written proceedings for the Twelfth British National Conference on Databases, BNCOD-12, held at the University of Surrey, Guildford, UK.

Despite the maturity of database technology, and its widespread acceptance as being central to software and system development, a number of recurring themes have presented themselves again at this conference, as they do also at other database conferences. These range from the implementation challenges of parallel and distributed systems through the problems of exploiting object-orientation in databases and capturing temporal data to the formal models and techniques which underpin the abstractions inherent in databases. Whilst it is true that all of the refereed technical papers have been submitted from academia, the issues which they address are critical for system development and implementation; hence, we would hope that there will be as much in this volume to interest the commercial practitioner as the academic researcher.

The environment in which databases operate has evolved, and this has in itself brought a range of new challenges. Our first invited speaker, *Richard Barker*, addresses one such issue – the management of open systems, including multi-vendor databases – which is particularly pertinent to ‘real’ database implementations, just as it is to most other applications.

When first developed, databases represented, on the whole, snapshots of the world at particular instants. Much effort has been expended investigating the modelling and representation of temporal data. Our first group of technical papers addresses this area: *Schwidorski and Saake* explore the underlying temporal logic which pertains to active databases, whilst *Soukeras and King* discuss an event-based approach to modelling the temporal dimension of data.

Databases rely fundamentally on underlying formalisms. Often, these formalisms become most apparent as the basis for new query languages or manipulation algorithms. *Chan and Trinder* present a query notation for object-oriented databases founded on list comprehension techniques, and *Alagic and Sunderraman* discuss the expressivity of various logic paradigms in the context of strongly typed object databases. In a third paper in this group, *Bagai and Sunderraman* discuss an extension to the relational model of data, and an associated algebra, which can be applied to the manipulation of general deductive databases.

As the availability of suitable hardware permits the development of massively parallel databases, a number of issues arise which are addressed in the next group of papers. *Jelly, Kerridge and Bates* present an approach to benchmarking parallel database machines, which have, of course, performance characteristics which are fundamentally different from those of traditional mainframes. *Burger and Thanisch* discuss modifications to the standard concept of transactions which would promote the achievement of massive concurrency in parallel systems. *McCarroll and Kerridge* explore the issue of semantic integrity and the derivation of integrity tests appropriate to parallel database machines.

Object-oriented databases are proving to be far from the simple idea which was first espoused, and they offer a wealth of challenges which are being addressed by a number of researchers. Here, we have three papers which discuss three different areas of concern: the impedance mismatch at the interface, efficient access to stored objects, and an ad-hoc query language for objects. The papers by *Paton, al-Qaimari and Doan, Kemp, Iriarte and Gray* and *Barclay and Kennedy*, respectively, discuss these three topics.

The final pair of technical papers explore the area of federations of heterogeneous distributed database systems and the requirement for interoperability. *Murphy and Grimson* describe the Jupiter system, a multidatabase 'layer' which provides interoperability language, and *Jeffery et al.* propose a 5-layer model of information content to aid the reconciliation of the data models for autonomous heterogeneous information sources.

Our second invited speaker, *Peter Gray*, draws the conference to a close by discussing the challenge of releasing the knowledge that is now widely available on the so-called 'information highway'. He suggests that a combination of object databases, client-server architectures and 'knowledge brokers' will be necessary for the full potential of knowledge networks to be realised.

Acknowledgements

The task of selecting 13 papers for presentation from the 47 submitted fell to the programme committee of some 23 members, to whom I am deeply grateful. In addition to the papers included here, others were recommended for poster presentation during the conference. Bill Olle was a tower of strength as chair of the committee, in the initial preparation of the papers for refereeing, the collation of referees' comments for the programme committee meeting and the drafting of the outline programme, as well as guiding the committee itself through the material to hand. As always, the programme committee has attempted to maintain a balance of interest between academia and commerce, and I am especially grateful to Bill for his leadership in this area.

I am particularly indebted to the guidance and support of the BNCOD steering committee, especially Alex Gray, its chair, who has responded to countless phone-calls and email messages, as well as offering much useful practical advice. In addition, the organisers of previous conferences – especially Mike Worboys and Peter Gray – have offered invaluable advice at several stages. The conference itself could not have been organised at Surrey without the long-suffering assistance of Helen Slee and Caroline Leishman, who handled most of the administrative details, and postgraduate students in the Department have also been a major source of effort behind the scenes.

Guildford, May 1993

David S. Bowers

Conference Committees

Programme Committee

T. W. Olle (T. William Olle Associates) - Chair
 T. J. Bourne (SIAM Ltd.)
 D. S. Bowers (University of Surrey)
 R. Cooper (University of Glasgow)
 B. Eaglestone (University of Bradford)
 D. J. L. Gradwell (Data Dictionary Systems Ltd.)
 P. M. D. Gray (University of Aberdeen)
 W. A. Gray (Univ College of Wales, Cardiff)
 J. Hughes (University of Ulster)
 M. Jackson (University of Wolverhampton)
 K. G. Jeffery (Rutherford Appleton Laboratory)
 J. Kennedy (Napier University)
 R. Lucas (Keylink Computers Ltd.)
 S. Malaika (IBM, Hursley)
 J. K. M. Moody (University of Cambridge)
 M. A. Newton (Open University)
 A. Poulouvassilis (Kings College, University of London)
 N. Revell (City University)
 C. Small (Birkbeck College, University of London)
 P. M. Stocker (University of East Anglia)
 R. Tagg (Independent Consultant)
 M. Worboys (University of Keele)
 M. H. Williams (Heriot-Watt University)

Steering Committee

W. A. Gray (Univ College of Wales, Cardiff) - Chair
 T. J. Bourne (SIAM Ltd.)
 D. S. Bowers (University of Surrey)
 P. M. D. Gray (University of Aberdeen)
 M. Jackson (University of Wolverhampton)
 M. Worboys (University of Keele)
 M. H. Williams (Heriot-Watt University)

Organising Committee

D. S. Bowers (University of Surrey)
 C. Leishman (University of Surrey)
 H. Slee (University of Surrey)

Contents

Invited Papers

- Managing open systems now that the "glasshouse" has gone 1
R. Barker (OpenVision International, Camberley, UK)
- Knowledge reuse through networks of large KBS 13
P.M.D. Gray (University of Aberdeen, UK)

Temporal Databases

- Expressing temporal behaviour with extended ECA rules 23
S. Schwiderski and G. Saake (University of Cambridge, UK)
- Temporal databases: an event-oriented approach 38
S. Soukeraas and P. J. H. King
(Birkbeck College, University of London, UK)

Formal Approaches

- Object comprehensions: a query notation for object-oriented databases 55
D. K. C. Chan and P. W. Trinder (University of Glasgow, UK)
- Expressivity of typed logic paradigms for object-oriented databases 73
S. Alagic and R. Sunderraman (Wichita State University, USA)
- Algebraic computation of the weak well-founded model for general deductive databases 90
R. Bagai and R. Sunderraman (Wichita State University, USA)

Parallel Databases

- Benchmarking parallel SQL database machines 105
I. Jelly (Sheffield Hallam University, UK),
J. Kerridge (University of Sheffield, UK) and
C. Bates (National Transputer Support Centre, Sheffield, UK)

Branching transactions: a transaction model for parallel databases <i>A. Burger and P. Thanisch (University of Edinburgh, UK)</i>	121
--	-----

A strategy for semantic integrity enforcement in a parallel database machine <i>N. McCarroll and J. Kerridge (University of Sheffield, UK)</i>	137
---	-----

Object Oriented Databases

On interface objects in object-oriented databases <i>N. W. Paton, G. al-Qaimari and K. Doan (Heriot-Watt University, UK)</i>	153
---	-----

Efficient access to FDM objects stored in a relational database <i>G. J. L. Kemp, J. J. Iriarte and P. M. D. Gray (University of Aberdeen, UK)</i>	170
---	-----

A conceptual language for querying object-oriented data <i>P. J. Barclay and J. B. Kennedy (Napier University, UK)</i>	187
---	-----

Distributed Databases

The Jupiter system: a prototype for multi-database inter-operability <i>J. P. Murphy (Dublin City University, Republic of Ireland) and J. B. Grimson (Trinity College, Dublin, Republic of Ireland)</i>	205
--	-----

A model for heterogeneous distributed database systems <i>K. G. Jeffery, L. Hutchinson, J. Kalmus, M. Wilson, W. Behrendt and C. Macnee (Rutherford Appleton Laboratory, UK)</i>	221
---	-----