

Advances in Databases

11th British National Conference on Dátabases, BNCOD 11 Keele, UK, July 7-9, 1993 Proceedings

Springer-Verlag

Berlin Heidelberg New York London Paris Tokyo Hong Kong Barcelona Budapest Series Editors

Gerhard Goos Universität Karlsruhe Postfach 6980 Vincenz-Priessnitz-Straße 1 D-76131 Karlsruhe, FRG Juris Hartmanis Cornell University Department of Computer Science 4130 Upson Hall Ithaca, NY 14853, USA

Volume Editors

Michael Frederick Worboys Anna Frances Grundy Department of Computer Science, Keele University Keele, North Staffordshire, ST5 5BG, UK

CR Subject Classification (1991): H.2

6204

ISBN 3-540-56921-9 Springer-Verlag Berlin Heidelberg New York ISBN 0-387-56921-9 Springer-Verlag New York Berlin Heidelberg

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 1993 Printed in Germany

Typesetting: Camera ready by author Printing and binding: Druckhaus Beltz, Hemsbach/Bergstr. 45/3140-543210 - Printed on acid-free paper

Foreword

The papers following constitute the written proceedings of the eleventh British National Conference on Databases, held at Keele University in the Midlands of England. As with many other current database conferences, a dominant theme is the provision of the means to enhance the capabilities of databases to handle information that has a rich semantic structure. Such provision is clearly based upon user requirements, since many applications demand it and system configurations call for it. A major research question is how to achieve such a semantic scale-up without sacrificing performance. There are currently two main paradigms within which it is possible to propose answers to this question, deduction-oriented and object-oriented. These two approaches are not disjoint; indeed, some would argue that the object-oriented approach is subsumed within the deductive logic-based approach. Both paradigms are well represented in this collection, with perhaps the balance swinging in the direction of the deductive approach, which is followed by hoth the invited papers.

It has been an unfortunate consequence of the admitted failure of the Japanese Fifth Generation project to develop an effective logic-based paradigm that all approaches based on logic are seen in some quarters as doomed to failure. Our first invited speaker takes a much more optimistic line. *Michael Freeston*, from the European Computer-Industry Research Centre (ECRC) at Munich has argued that ECRC has achieved more than both the Japanese and American initiatives in this area and that logic-based information systems are far from dead. His paper describes work to provide efficient indexes to support both deductive and object-oriented paradigms.

Both deductive and active databases are based upon rules, yet unification of these two types has not yet been achieved. Our second invited speaker, *Carlo* Zaniolo from the University of California at Los Angeles, discusses the technical problems at the heart of such a unification. He outlines a semantic framework that results in a rule-based language, presently being developed at UCLA, with the potential to integrate deductive and active databases.

Three of the papers following address the theme of semantic enhancement in the context of specific applications. Goble and Glowinski discuss the handling of integrity constraints in a medical application, Fernandes et al. test a generic unification of deductive and object-oriented approaches with a geographic application, and Eaglestone et al. consider how recent research on database support for design activities can extend beyond support for standard engineering and software design to the artistic design process.

Two papers discuss distributed configurations. Castellanos considers ways to enrich existing models in constituent databases of a distributed system so as to provide integrated access to the federation. Cohen and Ringwood use a logicbased language that supports concurrency to integrate heterogeneous nodes.

Four papers deal directly with the object-oriented approach to databases. Barclay and Kennedy incorporate the established database concept of the 'view' into their object-oriented world. Hohenstein and Odberg use an extended version of the entity-relationship model to enrich the semantics of databases and provide the basis for a uniform C++ interface to a variety of DBMS. Carnduff and Gray consider how function materialisation in object-oriented databases can enhance their use for the engineering design process. Tagg and Liew provide an excellent summary of the current state of object-oriented DBMS.

Returning to the deductive theme, Sunderraman constructs an extension to definite deductive databases that allows for conditional facts. Model-theoretic and fixpoint semantics are provided for this extension. Singleton and Brereton consider some of the practicalities in implementing logic features using a relational database.

There are two other papers in this collection which do not fit so easily into our general theme, yet are very much in the mainstream of current database research. *McBrien* discusses some of the algorithms required for querying and updating an historical relational database. The discussion provides information about what is required of a relational system to support an historical information system. *Reddi's* paper makes a contribution to the theory of functional databases by showing how work on optimisation strategies for integrity constraint enforcement in deductive databases can be transferred to the functional database approach.

Acknowledgements

I am very grateful to the programme committee, whose members are listed on a later page. The programme committee was chaired quietly and effectively by John Hughes at a meeting in February. We selected 13 papers for presentation at the eleventh BNCOD conference from a total of 36 full papers submitted. Other submissions were recommended for presentation in poster form. Each full paper was referred by at least three members of the programme committee.

I would like to acknowledge the support of the BNCOD steering committee, especially its chair, Alex Gray. Last year's conference organiser, Peter Gray, offered much useful practical advice. At Keele, I am grateful for the help of the conference organising committee, without whose assistance this conference could not have taken place. Frances Grundy was able to contribute the benefit of her experience of organising BNCOD at Keele in 1985. My postgraduate students have done a great deal of work behind the scenes. In particular, Priyantha Jayawardena shared many hours with me in routine tasks. Kendal Allen, the conference administrator, provided a friendly and efficient interface with delegates.

Keele, May 1993

Michael F. Worboys

Conference Committees

Programme Committee

J Hughes (University of Ulster) – Chair J Bocca (University of Birmingham) TJ Bourne (SIAM Limited) R Cooper (University of Glasgow) SM Deen (Keele University) B Eaglestone (University of Bradford) WA Gray (University of Wales, Cardiff) AF Grundy (Keele University) KG Jeffery (SERC) G Kemp (University of Aberdeen) JB Kennedy (Napier, Edinburgh) RJ Lucas (Keylink Computers Limited) Z Kemp (University of Kent) A Poulovassilis (King's College, London University) NW Paton (Heriot Watt University) R Tagg (Independent Consultant) GCH Sharman (IBM, Hursley) C Small (Birkbeck College, London University) MF Worboys (Keele University)

Steering Committee

WA Gray (University of Wales, Cardiff) - Chair
PMD Gray (University of Aberdeen)
MS Jackson (Wolverhampton University)
MH Williams (Heriot Watt University)
MF Worboys (Keele University)

Organising Committee

K Allen (Keele University) AF Grundy (Keele University) P Jayawardena (Keele University) MF Worboys (Keele University)

BIBLIOTHEQUE DU CERIST

Contents

Invited Papers

Technical Papers

Semantic Constraints in a Medical Information System
A Methodology for Semantically Enriching Interoperable Databases
Distributed Databases Tied with StrIng
Viewing Objects
Function Materialization Through Object Versioning in Object-Oriented Databases 111 T. W. Carnduff, W. A. Gray (University of Wales College of Cardiff, UK)
A C++ Database Interface Based on the Entity-Relationship Approach 129 Uwe Hohenstein (Siemens AG, Munich, Germany), Erik Odberg (Norwegian Institute of Technology, Trondheim, Norway)
Object-Oriented Database Methodology – State of the Art
Deductive Databases with Conditional Facts
A Deductive Object-Oriented Database for Data Intensive Application Development
Alvaro A. A. Fernandes, Maria L. Barja, Norman W. Paton, M. Howard Williams (Heriot-Watt University, UK)

Storage and Retrieval of First-Order Terms Using a Relational Database ... 199 Paul Singleton, O. Pearl Breseton (Keele University, UK)

Begriffsverzeichnis: a Concept Index

Michael Freeston

ECRC, Arabellastrasse 17, DW-8000 Munich 81, Germany

Abstract. This paper describes a generalised technique for the indexing of data structures of different complexities. It represents the culmination of a long research effort to devise efficient and well-behaved indexing for the lowlevel support of persistent programming languages: more specifically, persistent logic programming and object-oriented programming environments, which the author sees as the future of database systems.

1 Introduction

This paper describes a generalised technique for the indexing of data structures of different complexities. It represents the culmination of a long research effort to devise efficient and well-behaved indexing for the low-level support of persistent programming languages: more specifically, persistent logic programming and objectoriented programming environments, which the author sees as the future of database systems.

Both programming paradigms are particularly amenable to persistent implementation. In the author's conception of such systems, the executed code locates the target of the next execution step - whether the target be method, object, rule or ground fact - by indirect reference via an index (i.e. search by structure and value), rather than by direct reference to a compiled-in address location. In persistent logic programming, location via an index acts as a clause pre-unification step. In persistent object-oriented programming, it provides the location of complex objects, and the late binding of method to object.

The relative inefficiency introduced by indirect addressing through an index has to be balanced against the advantage of being able to apply all the conventional DBMS operations to objects or clauses, thereby extending the power of the programming language to the database system, while eliminating the 'impedance mismatch' which exists between every programming language and embedded database language. In deductive systems, the loose coupling between a logic programming language and a (relational) DBMS is transformed into the basis of a true, fully integrated, deductive database system. In the object-oriented paradigm, it is the author's view that, whatever arguments there may be about the precise properties which an object-oriented database system should possess, it is nothing if it is not a persistent object-oriented programming language.

Efficient, generalised indexing of complex structures is thus an essential prerequisite, if not the key, to viable persistent programming. But it has a more fundamental significance: it is 120 years since Frege devised a notation (Begriffsschrift) for the representation of any arbitrary concept. Since we reason by matching the structure of one concept against another, efficient reasoning on large bodies of knowledge in