

Lecture Notes in Computer Science

Edited by G. Goos and J. Hartmanis

466

A. Blaser (Ed.)

Database Systems of the 90s

International Symposium

Müggelsee, Berlin, FRG, November 5–7, 1990

Proceedings



Springer-Verlag

Berlin Heidelberg New York London

Paris Tokyo Hong Kong Barcelona

Editorial Board

D. Barstow W. Brauer P. Brinch Hansen D. Gries D. Luckham
C. Moler A. Pnueli G. Seegmüller J. Stoer N. Wirth

Volume Editor

Albrecht Blaser

IBM Deutschland GmbH

Pascalstr. 100, W-7000 Stuttgart 80, FRG

BIBLIOTHEQUE DU CERIST

5437

CR Subject Classification (1987): H.2, H.1.2

ISBN 3-540-53397-4 Springer-Verlag Berlin Heidelberg New York

ISBN 0-387-53397-4 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 other ways, and storage in data banks. Duplication of this publication or parts thereof is only permitted under the provisions of the German Copyright Law of September 9, 1965, in its current version, and a copyright fee must always be paid. Violations fall under the prosecution act of the German Copyright Law.

© Springer-Verlag Berlin Heidelberg 1990

Printed in Germany

Printing and binding: Druckhaus Beltz, Hemsbach/Bergstr.

2145/3140-543210 – Printed on acid-free paper

PREFACE

The papers of these proceedings were presented at the International Symposium of IBM Germany on Database Systems of the 90s, held November 5-7, 1990 at Berlin.

The aim of IBM Germany's scientific symposium series is to strengthen and improve the scientific communication between academia and industry by covering subjects related to the fields of computer science and application, both from an academic and from an industry point of view. At IBM Germany we are all extremely happy that the political changes in the Eastern World allowed us to choose a conference site in East Germany for the first time in the long history of this series. Doing so - and inviting some 50 East German scientists to attend - we wanted again to demonstrate that we are determined to fully integrate this part of Germany and our fellow scientists working there into our scientific activities.

The evolution of database system technology in the 80s was characterized by

- Relational database products penetrating and driving the computer market on all levels - personal computers, powerful workstations and mainframes - with relatively stable functionality and user interfaces, and product evolution concentrating on performance and distribution
- A large portion of database research concentrating on enhancing the modeling power of the (relational) data model towards more complex objects than just those easily representable by "flat tuples and tables", to cope with so-called non-standard applications (in office, engineering, and science); from there, identifying and tackling related problems such as transaction management, (time) versioning, extensibility, workstation-server cooperative processing, etc.
- Rich research on the more classical architectural, algorithmic and performance issues of relational database systems with topics like high transaction load, distribution, hardware support, etc.

At the change of the decade, the question is obvious: What will be the dominating driving forces, activities and evolutions in the database systems scene of the 90s, in research, ad-tech, products, and applications? This, coupled with the observation that it has been already 15 years that IBM Germany devoted one of its scientific symposia to database systems, motivated us to organize this conference. The scope of the conference and its 14 invited papers address

- requirements and demands,
- user aspects,
- architectural and system aspects for novel applications,
- and system and implementation aspects

of database systems. This cannot - of course - exhaust the myriad of interesting and important issues facing the database community at present. We had to be selective, and the preference was naturally influenced by the database research which has been conducted during the past 15 years at IBM Germany's Heidelberg Scientific Center with quite some effect on the evolution of relational database technology for non-standard applications.

The keynote paper of the conference gives a historical account of the evolution of database technology for the classical application area of enterprise information systems (with heavy demands for performance under high transaction load) and projects the demands and the evolution into the 90s.

Session 2 deals with the demands and driving forces of "non-standard" applications in office, engineering, and science, and very specifically with those of multi-media complex objects. Another driving force - particularly for database system products, but to some extent also for database research - comes from standardization efforts.

Session 3 deals with two classes of database system users and their language requirements: application programmers, who need a combination of programming and data definition/manipulation languages, and end users, who need an on-line, interactive query language to formulate and execute ad-hoc queries.

The selection of papers in session 4, dealing with database system and architecture concepts as requested by novel applications, is perhaps the one most biased

by the research interests of the Heidelberg Scientific Center. Evolution of data models, concepts and perspectives of object oriented database systems, deductive database rule languages, extensibility, and data replication (e.g., for performance and reliability, particularly in distributed and cooperative database systems) are felt to be of great importance for the future of database systems and applications, and they are more or less touched by our own research.

Session 5 on system and implementation aspects covers just three of the many topics which we felt will have significant impact on database systems in the 90s. Distributed and cooperative database management is dictated by many classical and non-standard applications. Also, performance and reliability will remain in the focus of usable systems (there are hardly any other computer applications having higher performance and reliability demands than enterprise information systems based on database systems). And finally, the impact of hardware on systems architecture, performance and user behaviour is one of the most obvious questions which has always been posed in the context of database systems, and always will.

The symposium concluded with a panel discussion, dwelling on the speculative question "Quo Vadis DBMS?".

All papers have been invited. And although the symposium is sponsored by IBM Germany, the selection of topics reflects the editor's personal judgement of their significance and not IBM's. Similarly, the contents of the papers express the authors' own personal opinions.

The authors were explicitly invited not to contribute original research papers but rather to give an account - in the light of their own research - on where database technology (products, research, applications) in their respective subject area stands at the change of the decade, what the present highlights of research are and - as a consequence - what is likely to happen in the 90s. I am very grateful that all authors spontaneously agreed with the idea of the conference and with the nature of the contributions, which sometimes forced them to be speculative - a behaviour which scientists usually don't like - and which caused for most of them much more effort and time to write the paper than if they had just reported about their latest research.

The symposium was sponsored by IBM Germany and organized by its Scientific Relations Department. I would like to express my gratitude and appreciation to the sponsor and organizers (in particular to Dr. H. Statz, Dr. W.-D. Oberhoff, and Mrs. I. Knödel), to the authors, lecturers, panelists, and session chairpersons, and to the many contributors within and outside of IBM (in particular to Dr. K. Küspert, the manager of the Advanced Information Management Research Department of the IBM Scientific Center at Heidelberg and to his research staff) who gave me their advice and assistance in preparing, organizing, and running the symposium, and also in reading and commenting on the papers.

As another "first", this symposium will be recorded on tape by EuroPACE, the European Program of Advanced Continuing Education at Paris. The tapes will be distributed via the EuroPACE Network or mail to the EuroPACE Member Organizations and to requesting universities in the former German Democratic Republic. My thanks go to the management and to the staff of EuroPACE (in particular to Tage E. Frisk, President and General Manager of EuroPACE, and to Sören Nipper, Project Manager) as well as to the camera team from Media-Design Stuttgart.

Last but not least my cordial thanks go to Springer-Verlag in Heidelberg (in particular to Dr. H. Wössner) for rather spontaneously agreeing to publish the proceedings in the Lecture Notes in Computer Science series, and for producing the volume in the obviously unavoidable last minute hurry, but nevertheless in the well known, outstanding Springer quality.

Heidelberg, October 1990

A. Blaser

CONTENTS

1. Keynote Paper

Database Role in Information Systems: The Evolution of Database Technology and its Impact on Enterprise Information Systems

Donald J. Haderle

1

2. Demands on Database Systems in the 90s

Future Database Technology: Driving Forces and Directions

Peter C. Lockemann, Alfons Kemper, Guido Moerkotte

15

Management of Multimedia Complex Objects in the '90s

Michel E. Adiba

34

Database Language Standards: Past, Present, and Future

Phil Shaw

55

3. User Aspects

Language Technology for Post-Relational Data Systems

Joachim W. Schmidt, Florian Matthes

81

Human Factors of Database Systems

Rainer Ginnich, Klaus Kunkel, Thomas Strothotte

115

4. Database System and Architecture Concepts for Novel Applications

Evolution of Data Models

Hans-Jörg Schek, Marc H. Scholl

135

Object-Oriented Database Systems: Concepts and Perspectives

Rainer Unland, Gunter Schlageter

154

Deductive Database Rule Languages: Analysis and Case Study

Georges Gardarin, Jerry Kiernan

198

Extensibility in the Starburst Experimental Database System

Bruce Lindsay, Laura Haas

217

Logical Foundation of Data Replication

Stefan Jablonski, Hartmut Wedekind

249

5. System and Implementation Aspects**Trends in Distributed and Cooperative Database Management***Klaus Küspert, Erhard Rahm***263****Performance and Reliability Issues in Future DBMSs***Andreas Reuter***294****The Impact of Hardware on Database Systems***Patricia G. Selinger***316**

Database Role in Information Systems: The Evolution of Database Technology and its Impact on Enterprise Information Systems

Donald J. Haderte

**IBM Corporation
555 Bailey Ave
San Jose, California 95141**

Abstract

During the 1960's and 1970's commercial database management systems focused on transaction and batch processing environments with an emphasis on minimizing computing resources. Relational database management systems appeared on the commercial scene in the 1980's and extended the range of applications amenable to database processing. These applications now include interactive environments as well as transaction and batch. Relational DBMSs emphasize function, ease of data access by users, and improved economics of application development at the expense of computing resources. This paper looks at how enterprise information systems have been affected by the evolution of database technology in the 1980's, emphasizing the role of relational database and forecasting expectations for the 1990's.

Pioneering Commercial Data Models

The pioneering commercial database management systems in the 1960's and early 1970's focused on the development of application models for managing concurrent read/write access to data with an emphasis on minimizing use of computing resources. Batch processing dominated the computing landscape and continues to play a major role today as an application model for work within information systems.

At the same time, techniques for application work management known as transaction management systems evolved. Transaction management systems provided the framework for formally specifying which application work (called a transaction) should be considered as an atomic set of operations whose changes to data must be manifested in the database if the transaction succeeds. If the transaction fails, none of the changes made by its operations can appear in the database. Transaction management systems provided an environment for on-line transaction processing, the execution of a few short, interactive database operations as a single transaction. The need for on-line transaction processing arose from the demand for information currency and real-time services which batch processing could not provide. The ensuing