UNIVERSITY OF ABERDEEN

ICOD-2

PROCEEDINGS

Second International Conference on Databases

BRITISH COMPUTER SOCIETY WORKSHOP SERIES

CHURCHILL COLLEGE, CAMBRIDGE, SEPTEMBER, 1983

Edited by S. M. DEEN and P. HAMMERSLEY

University of Aberdeen

The British Computer Society

ICOD-2

PROCEEDINGS

Second International Conference On Databases

Churchill College, Cambridge
August 30 - September 3 1983

Edited by

S. M. Deen
University of Aberdeen

and

P. Hammersley
The Computer Journal

Published by
WILEY HEYDEN LTD
on behalf of
THE BRITISH COMPUTER SOCIETY

British Library Cataloguing in Publication Data

International Conference on Databases *(2nd: 1982 Cambridge)*
ICOD-2.—(The British Computer Society workshop series)
1. Database management—Congresses
I. Title II. Deen, S. M.
III. Hammersley, P. IV. Series
001.64'42 QA76.9.D3
ISBN 0 471 90309 4

Printed in Great Britain

ICOD-2

Second
International Conference
on
Databases

THE BRITISH COMPUTER SOCIETY WORKSHOP SERIES

Editor: P. Hammersley

The BCS Workshop Series aims to report developments of an advanced technical standard undertaken by members of The British Computer Society through the Society's study groups and conference organisation. The Series should be compulsive reading for all whose work or interest involves computing technology and for both undergraduate and post-graduate students. Volumes in this Series will mirror the quality of papers published in the BCS's technical periodical The Computer Journal and range widely across topics in computer hardware, software, applications and management.

Some current titles:

Data Bases: Proceedings of the International Conference 1980
Eds. S. M. Deen and P. Hammersley

Minis, Micros and Terminals for Libraries and Information Services
Ed. Alan Gilchrist

Information Technology for the Eighties BCS '81
Ed. R. D. Parslow

ICOD-2

Second
International Conference on Databases
Cambridge, UK, Aug 30 - Sep 3, 1983

ORGANISERS

The British Computer Society
University of Aberdeen
University of Cambridge
Université de Paris-Sud

SPONSORS

AFCET (France)
AICA (Italy)
European Research Office of the US Army
GI (FRG)
International Computers Limited
IEEE (USA)
INRIA (France)
Shell International Petroleum Company Limited
In cooperation with ACM (USA)

ICOD-2

SECOND
INTERNATIONAL CONFERENCE ON DATABASES
Cambridge, UK, Aug 30 - Sep 3, 1983

PROGRAMME COMMITTEE

Chairman: S. M. Deen *Universitry of Aberdeen, UK*

P. Buneman *University of Pennsylvania, USA*
G. Bracchi *Politecnico di Milano, Italy*
E. Falkenberg *University of Queensland, Australia*
G. Gadarin *INRIA, France*
E. Gelenbe *Université de Paris-Sud, France*
P. Hammersley *The Computer Journal, UK*
P. J. H. King *Birkbeck College, London, UK*
W. Litwin *INRIA, France*
F. A. Manola *Computer Corporation of America, USA*
E. J. Neuhold *Universität Stuttgart, FRG*
J. S. Riordon *Carleton University, Canada*
F. A. Schreiber *Universita Degli Studi di Parma, Italy*
A. Sølvberg *University of Trondheim, Norway*
P. M. Stocker *University of East Anglia, UK*
R. P. van de Riet *Vrije University, The Netherlands*
R. Williams *IBM, San Jose, USA*

Secretariat: G. O. Ofori-Dwumfuo *University of Aberdeen, UK*

WORKSHOP SUBCOMMITTEE

Chairman: E. Gelenbe
G. Gadarin, W. Litwin, F. A. Manola, R. Williams

ORGANISING COMMITTEE

Chairman: S. M. Deen *University of Aberdeen, UK*
Secretary: P. Hammersley *The Computer Journal, UK*
D. F. Hartley *University of Cambridge*
J. K. M. Moody *University of Cambridge*
B. Westwood *University of Cambridge*

LIST OF REFEREES (ICOD-2)

R.R. AMIN	Britain	G. M. LOHMAN	USA
P. M. G. APERS	Netherlands	R. LORIE	USA
P. ATZENI	Italy	J. MADELAINE	France
D. A. BELL	Ireland	S. A. MAHMOUD	Canada
P. BERNADAT	France	F. A. MANOLA	USA
P. BOUCHET	France	G. MARTELLO	Italy
D. S. BOWER	Britain	N. MARTIN	Britain
G. BRACCHI	Italy	I. NARANG	USA
P. BUNEMAN	USA	E. NEUHOLD	Germany
U. BUSSOLATI	Italy	G. O. OFORI-DWUMFUO	Britain
S. CERI	Italy	V. de OUTONETHIS	Italy
A. CHAN	USA	G. PELAGATTI	Italy
J. DAVIDSON	USA	B. PERMICI	Italy
U. DAYAL	USA	A. REUTER	Germany
S. M. DEEN	Britain	J. S. RIORDON	Canada
J. A. EDGAR	Britain	O. RISNES	Norway
E. FALKENBERG	Australia	G. SCHLAGETER	Germany

P. FAUDEWAY	France	J. W. SCHMIDT	Germany
S. FINKELSTEIN	USA	F. SCHREIBER	Italy
J. J. FLORENTIN	Britain	K. SHOENS	USA
S. FOX	USA	J. M. SMITH	USA
G. GARDARIN	France	A. SOLVBERG	Norway
E. GALENBE	France	W. STANISZKIS	Italy
C. GHEZZI	Italy	F. STEYER	Germany
G. GOTTLOB	Italy	P. M. STOCKER	Britain
L. HAAS	USA	M. C. TAYLOR	Britain
P. HAMMERSLEY	Britain	R. TAYLOR	USA
G. JOMIER	France	P. TIBERIO	Italy
W. de JONGE	Netherlands	K. C. TOTH	Canada
M. JOUVE	France	R. P. VAN de RIET	Netherlands
J. JUGLIS	Britain	P. VALDURIEZ	France
M. L. KERSTEN	Netherlands	A. VASHISHTA	Britain
P. J. H. KING	Britain	B. WALTER	Germany
C. H. KUNG	Norway	R. A. WILLIAMS	USA
B. LINDSAY	USA	C. M. WOODSIDE	Canada
W. LITWIN	France	R. YASOLI	Germany

PREFACE

This conference is a follow-up to the first International Conference On Databases (ICOD-1) held at Aberdeen, Scotland in July 1980, and forms part of the Silver Jubilee Celebrations of The British Computer Society. It is being organised jointly by The British Computer Society and the Universities of Aberdeen and Cambridge. An additional feature of this conference is a special workshop on New Applications of Databases, organised by Informatique des Systemes Experimentaux et leur Modelisation, Universite de Paris-Sud.

The objective of the ICOD-2 is to stimulate database research in Europe by providing a forum where researchers, experts and other interested parties from all over the world can discuss and debate issues of interest. Special emphasis has been placed on new ideas, future directions and application-oriented research.

The main components of ICOD-2 are:

Main Conference (30 August - 2 September, 1983)

Invited Papers
Contributed Papers
Panel Discussions
Posters

Workshop (2 - 3 September, 1983)

New Applications of Databases

We have three invited papers, covering some topical areas of the database technology. Irving L. Traiger (IBM, San Jose, USA) reviews some of the mismatches between the operating system and DBMS, and discusses the interactions of the system components, advocating a more coherent approach for improved performance. Frank A. Manola and Alain Pirrote (Computer Corporation of America, USA and Phillips Research Laboratory, Belgium) investigate the problems of developing common standards for Codasyl, relational and hierarchical models and suggest a unified framework for a family of standards. Herbert Weber (Universitat Bremen, FRG.) presents some basic concepts of an object-oriented language suitable for the specification and design of distributed databases.

There were about 80 papers submitted to the conference, from which 16 papers, referred to as the contributed papers, have been selected for formal presentation, each paper being previously reviewed by 4 independent referees. These contributed papers are divided into the following six sessions:

Database Machines (2 papers)
Query Facilities (3 papers)
Conceptual Modeling (3 papers)
New Functionalities (3 papers)
Distributed Database Techniques (3 papers)
Performance Evaluation (2 papers)

We describe these papers briefly below:

Database Machines

L. Bic, and R. L. Hartmann (University of California, Irvine, USA) propose a database model - based on the principles of data flow system - for implementation on a database machine; the model has a number of interesting features with potential for further development. C. H. Pygott (Royal Signals and Radar Establishment, Ministry of Defence, U.K) describes a micro-programmable machine which is being developed at the British Ministry of Defence in order to process efficiently a wide variety of relational queries from a small number of relations.

Query Facilities

C. J. Date (IBM, San Jose, USA) expounds the implementation issues of the outer join operation and proposes a technique for its implementaion in full generality within a SQL environment. The next paper given by A. Motro (University of Southern California, Los Angeles, USA) describes a mechanism for integrating schemas of several independent databases into a common form, such as those required for global queries in a distributed database. The final paper of this session is contributed by S. Christodoulakis (University of Toronto, Canada) who presents a technique for batching queries - to improve performance in a concurrent environment - by means of access tables yielding qualifying records.

Conceptual Modelling

T. L. Anderson (Burroughs Corporation, USA) presents a data model, based on the concepts of events and processes, in order to support the impact of time dimension on both information and constraints. E. Wong (Universty of California, Berkeley, USA) proposes a mechanism to support abstract domain types on QUEL through virtual operations on Views, which provide, among others, geometric and text processing capabilities. D. Stemple, K Ramamritham, S. Vinter and T. Sheard (University of Massachusetts, USA) present a model which uses the operating system to support distributed databases as abstract data types, simplifying as a result some of the security and operational problems of databases.

New Functionality

M. A. Melkanoff and Q Chen (University of California, Los Angeles, USA) describe an experimental database which provides dynamic user facilities, such as the specification of integrity constraints or the monitoring of the database performance, by DM commands. E. Gelenbe

(Universite Paris-Sud and Centre Mondial, France) draws our attention to the need to provide facilities to represent 'summarised' information in a required format, from a detailed set of stored information; he also suggests some new concepts. J. B. Crampes, C. Y. Chrisment and G. Zurfluh (I.U.T. Informatique, France) present some results from a major database research project, purported to develop a data model and appropriate languages for ordinary, text, image and graphics data.

Distributed Databases

The first paper of this session, given by R. Unland, U. Praedel and G. Schlageter (University of Hagen, FRG.), presents a new version of optimistic concurrency controls based on a snapshot validation scheme for improved performance. The second paper by W. S. Luk and L. Luk (Simon Fraser University, Canada) describes an algorithm which can perform any given program for semi-join operation into a 'better' one which minimises the intersite communications cost. The final paper is presented by E. Bertino, C. Meghini, C. Pelagatti and C. Thanos (Politecnico di Milano and IEI-CNR, Italy) and it describes an update mechanism for a distributed database supporting distributed fragments and location transparency.

Performance Evaluation

W. Staniszkis, P. Rullo, M. Gaudioso and S. Orlando (CRAI, Italy) discuss a technique to improve the performance of Codasyl DM commands by holding distributed information with the data. The next paper, which is also the last paper of the main conference, is presented by N. H. Phuc, M. Becker and P. Sevray (CII-Honeywell Bull and Universite Paris VI, France) who compare the performances of Balanced trees with Prefix binary trees, showing that in the case of direct access the latter can be more efficient.

In addition to these papers, we also have three panel sessions with the following titles:

Personal Databases - the first ten years?
Distributed Databases - the next ten years?
Database Machines - the last ten years?

The panel titles are meant to be provocative rather than definitive. A poster session is planned for informal presentation of some recent developments in the general database area as posters.

This volume contains the three invited papers and the sixteen contributed papers in order of their presentation. It does not include any material from the Workshop, which is intended to be published as a separate volume at a future date. Unlike the main conference, the papers for the workshop have been selected on the basis of extended abstracts rather than full papers, and they are grouped into the following five presentation sessions:

Image Databases (2 papers)
Text Oriented Systems (3 papers)
CAD/CAM Applications (2 papers)
AI in Database Management (2 papers)
High Level Interfaces (3 papers)

with a sixth session under the title: "Databases - where next?" for an open discussion.

Finally I wish to thank all those who helped us in this conference, in particular to our organisers, sponsors, members of the organising and programme committees and referees - without whose help there would not have been a conference. My special thanks to the members of the PRECI database group at Aberdeen, the staff of the University of Aberdeen, The British Computer Society, Churchill College (Cambridge) and Middlesex Polytechnic, for their continuous assistance so essential for the success of a conference.

S. M. Deen
Conference Chairman

CONTENTS

Chapter 1

TRENDS IN SYSTEMS ASPECTS OF DATABASE MANAGEMENT

Irving L. Traiger

IBM Research Laboratory, Dept. K55, Bldg. 281, 5600 Cottle Road, San Jose, California 95193, USA

Considerable progress has been made on the "systems" side of database management, particularly in the area of serializability theory, and various mechanisms such as concurrency control, recovery, and process structure. Even though operating systems provide many of these mechanisms, they typically do not have the right performance or function to support database requirements. More recently, interest has grown in the general area of hardware and software systems, and how the various components should interact. It would be attractive to introduce database concepts into the operating system, so that a more coherent approach can be taken to global resource issues like concurrency control, process management, memory management, and recovery. In addition, the use of multiple, interconnected processors can offer increased capacity and availability for a database environment. This paper reviews some of the mismatches between database systems and operating systems, and discusses how a closer coupling could be obtained. It then discusses some multi-processor configurations, and the particular hardware and software environment that is being used at IBM San Jose Research Laboratory to explore system-wide issues.

1. INTRODUCTION

It is tempting to play the role of futurist for database systems, and consider what great new leaps will occur. To some extent, the work in relational systems has already given us a leap forward. These systems provide gains in usability and flexibility that are becoming increasingly important, as a wider variety of people begin to use database systems. Some researchers are trying to stretch those gains even further, through natural language and expert system techniques. Others are looking at unconventional machine organizations which are tuned to high level relational operators, and can work efficiently on large relations.

The emphasis in this paper is quite different, and is based on a number of assumptions:

1. Much of the database workload will continue to be simple, repetitive access. Examples include banking systems, reservation systems, order entry, and tracking of process steps in manufacturing. Even on the System R project [1] and its derivatives, we found ourselves estimating instruction counts and I/O activities with two canonical transactions, one from banking and the other from manufacturing.

2. Very high levels of concurrency will be required, involving large numbers of end user terminals. These terminals will be relatively intelligent, with local database capability, and the additional capability to invoke database functions on central computers. There are many computer installations today with large, interconnected terminal networks, and high projected growth. In the future, home computers and extended video cable capabilities will offer the potential for much more interactive database activity.

3. Large, central site computing facilities will continue to be important, especially for the support of database systems. Networks of departmental computers or personal workstations will not be suitable for the types of workloads described above. Much of the data will continue to be inter-related, and the response time and processing costs of distributed database protocols will be excessive for high volume applications.

4. The required processing power will exceed the capability of any conventional, uni-processor organization. Even today, many of the fastest conventional processors are internally organized as two-way or even four-way multi-processors. From a software perspective, it would be preferable to wait for quantum leaps in hardware, such as Josephson junction technology. However, even these enhancements may not provide sufficient capacity in the long term.

5. Availability of these large, central database systems will be increasingly important. As the size and business impact of these systems increase, hardware and software techniques will be needed to mask internal errors and limit the scope of failure. Clearly a general power failure is outside the purview of this work. but systems will be structured to speed up their restart even after a massive power failure. Prioritized and partial restart techniques, along with judicious use of non-volatile memories for checkpoint information, could greatly improve the situation.

6. Large central facilities will be made up of multiple, cooperating processors, each of which has a relatively conventional machine architecture. The use of multiple processors is really a corollary to the assumptions about capacity and availability. However, the way in which these processors should be organized and used to share data is more difficult to predict. In terms of machine architecture, there are certainly interesting areas to explore, such as reduced instruction set techniques to implement conventional programming languages more efficiently than today