

P. M. D. Gray R. J. Lucas (Eds.)

Cc 01-618



مكتبة سيرست  
Bibliothèque CERIST

# Advanced Database Systems

10th British National Conference on Databases,  
BNCOD 10  
Aberdeen, Scotland, July 6-8, 1992  
Proceedings

BIBLIOTHEQUE DU CERIST

**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  
W-7500 Karlsruhe, FRG

Juris Hartmanis  
Department of Computer Science  
Cornell University  
5149 Upson Hall  
Ithaca, NY 14853, USA

## Volume Editors

Peter M. D. Gray  
Rob J. Lucas  
University of Aberdeen, Department of Computing Science, King's College  
Aberdeen, AB9 2UB, Scotland, UK

6242

CR Subject Classification (1991): H.2

ISBN 3-540-55693-1 Springer-Verlag Berlin Heidelberg New York  
ISBN 0-387-55693-1 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 1992  
Printed in Germany

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

## Preface

The theme of this book is the potential of new *Advanced Database Systems*. When the first of this series of conferences was held in Aberdeen, back in 1980, the potential of relational databases was clear for all to see, but commercial systems were still a few years off. At that time CODASYL and IMS reigned supreme, whilst the PC did not even exist! During the 80's, relational databases have become an industry standard, but their shortcomings have also become more obvious, many of which were foreseen. In particular, there was a serious shortcoming in data modelling constructs in order to capture complex CAD applications, and a difficulty in extending the model to include new datatypes. This led to an interest in *Object-Oriented Databases*, addressed by three papers.

Following these, we have a paper by an invited speaker *David Gradwell* on how to extend entity-relationship modelling techniques, widely used on relational databases, so that they cope with object-oriented databases. This is an important requirement for their widespread adoption.

Many people have questioned whether relational databases could meet the performance standards of finely-tuned IMS databases. Fortunately there have been developments in the UK in the use of networks of parallel processors. These have now come of age, and we have thought it right to include amongst research papers two studies of state of the art relational databases running on *Parallel Hardware* - NCR/Teradata and Meiko. These show great promise for the future, particularly with regard to the scalability of the technology and the use of faster processors. Together with these is a research paper on the use of transputer networks.

The availability of parallel hardware leads nicely into the question of the use of Distributed Systems. Our second invited speaker *Dr Michael Brodie* from GTE Labs, USA, discusses the promises for these systems, where application programs and databases cooperate together over a network. But this leads on to the awkward problem of *Legacy Systems* which are crucial to an organisation, but very hard to migrate from. How are we going to change over and get the advantages of Advanced Systems?!

One of the big problems of early database systems was the so-called impedance mismatch between them and programming languages, both in the range of data types that could be stored, and in the handling of large sets. This has led to a range of implementations coupling various new programming languages to databases - Functional, Persistent and Logic Programming. It has also led to more advanced *Conceptual Data Models* which are implemented in these languages, as described by three papers.

One particular feature of advanced data models is the need, long foreseen by Nijssen, to capture as much of the semantics of the data, including constraints and active rules, in the database. *Logic Programming* systems have been very useful for this purpose, and our final section discuss systems that use it for active rules, program transformations and constraints.

I am particularly grateful to the program committee of 16 people who picked the selection of 12 papers from 36 submitted. Although there was a wide variety of topics, it is interesting that these themes emerged so strongly for presentation at the 10th BNCOD conference. I am also very grateful for the industrial members of the program committee, who guided us in the selection of industrial papers, using slightly different criteria from the academic papers. Thus we hope that people working in industrial software development, and engineering design will find much to interest them in this book.

## Acknowledgements

Finally, may I say that these proceedings represent the strength of a vigorous database community in Britain, without whose support the annual conference could not take place. I am particularly grateful for advice and support from Alex Gray, as steering committee chairman, and from Elizabeth Oxborrow and Mike Jackson who ran earlier conferences. Nearer home, I am very grateful to my co-editor Rob Lucas, to my research students, who have helped organise the conference and to the departmental secretaries Irene Kavanagh and Katie Blanchard who coped with much of the day to day work and are the unsung heroes!

Most of the work in the UK has been supported over the years by the Science and Engineering Research Council, to whom we are all grateful, and it is to be hoped that they will see the fruits of this research in these proceedings as an encouragement to them to support the proposed special initiative in Advanced Databases and Large Knowledge Bases from 1992.

April 1992

Peter M D Gray  
Conference Chairman

## Conference Committees

### Program Committee

P.M.D. Gray (University of Aberdeen) - Chair  
J. Bocca (University of Birmingham)  
T.J. Bourne (SIAM Limited)  
W.A. Gray (University of Wales, Cardiff)  
A.F. Grundy (University of Keele)  
M.S. Jackson (Wolverhampton Polytechnic)  
K.G. Jeffery (SERC)  
R.G. Johnson (Birkbeck College, London University)  
J.B. Kennedy (Napier, Edinburgh)  
R.J. Lucas (Keylink Computers Limited)  
J.K.M. Moody (University of Cambridge)  
M. Norrie (University of Glasgow)  
E.A. Oxborrow (University of Kent)  
N.W. Paton (Heriot-Watt, Edinburgh)  
G.C.H. Sharman (IBM, Hursley)  
M.J.R. Shave (University of Liverpool)

### Steering Committee

W.A. Gray (University of Wales) - Chair  
P.M.D. Gray (University of Aberdeen)  
A.F. Grundy (University of Keele)  
M.S. Jackson (Wolverhampton Polytechnic)  
E.A. Oxborrow (University of Kent)  
M.H. Williams (Heriot Watt University)

### Organising Committee

R.J. Lucas (Keylink Computers Limited)  
G.J.L. Kemp (University of Aberdeen)  
P.M.D. Gray (University of Aberdeen)  
D. Nikodem (University of Aberdeen)  
S.M. Embury (University of Aberdeen)  
S. Leishman (University of Aberdeen)



## Contents

### Invited Papers

The Promise of Distributed Computing and the Challenges of Legacy Systems . . . . .	1
<i>M.L. Brodie (GTE Labs, Boston, MA, USA)</i>	
Object-Oriented Requirements Capture and Analysis – The Orca Project . . . . .	29
<i>D.J.I. Gradwell (Data Dictionary Systems Limited, Camberley, UK)</i>	

### Object-Oriented Databases

A Model for Versioning of Classes in Object-Oriented Databases. . . . .	42
<i>S.R. Monk, I. Sommerville (Lancaster University, UK)</i>	
A Storage Manager for the Hypermode Model . . . . .	59
<i>E. Tuv, A. Poutovassilis, M. Levene (University College, London, UK)</i>	
Views and Formal Implementation in Three-Level Schema Architecture for Dynamic Objects . . . . .	78
<i>G. Saake, R. Jungclauss (Technical University Braunschweig, FRG)</i>	

### Parallel Implementations and Industrial Systems

The Meiko Computing Surface: A Parallel and Scalable Open Systems Platform for Oracle . . . . .	96
<i>A. Holman (Meiko Limited, Bristol, UK)</i>	
A Study of a Parallel Database Machine and its Performance – The NCR/Teradata DBC/1012 . . . . .	115
<i>J. Page (NCR/Teradata Limited, Chertsey, UK)</i>	
Control of a Large Massively Parallel Database Machine Using SQL Catalogue Extensions, and a DSDL in Preference to an Operating System . . . . .	138
<i>M. Unwalla, J. Kerridge (The University of Sheffield, UK)</i>	

## Non-Relational Data Models

Integration of Modal Logic and the Functional Data Model . . . . .	156
<i>D.R. Sutton, P.J.H. King (Birkbeck College, University of London, UK)</i>	
The Raleigh Activity Model: Integrating Versions, Concurrency, and Access Control . . . . .	175
<i>M.H. Kay, P.J. Rivett, T.J. Walters (International Computers Limited, Reading, UK)</i>	
A Graphical Data Modelling Program with Constraint Specification and Management . . . . .	192
<i>R. Cooper, Z. Qin (University of Glasgow, UK)</i>	

## Logic Programming and Databases

Association Merging in a Schema Meta-Integration System for a Heterogeneous Object-Oriented Database Environment . . . . .	209
<i>M.A. Qutaishat, N.J. Fiddian, W.A. Gray (Univ. of Wales, College of Cardiff, UK)</i>	
Generating Active Rules from High-Level Specifications . . . . .	227
<i>O. Díaz (Univ. of Aberdeen, UK, and Univ. of País Vasco, San Sebastián, Spain), S.M. Embury (University of Aberdeen, UK)</i>	
Employing Integrity Constraints for Query Modification and Intensional Answer Generation in Multi-Database Systems . . . . .	244
<i>M.M. Fonkam, W.A. Gray (University of Wales, College of Cardiff, UK)</i>	



# THE PROMISE OF DISTRIBUTED COMPUTING AND THE CHALLENGES OF LEGACY SYSTEMS

Michael L. Brodie  
Intelligent Database Systems  
GTE Laboratories Incorporated  
brodie@gte.com

## ABSTRACT

The imminent combination of computing and telecommunications is leading to a compelling vision of world-wide computing. The vision is described in terms of next generation computing architectures, called Enterprise Information Architectures, and next generation information systems, called Intelligent and Cooperative Information Systems. Basic research directions and challenges are described as generalizations of corresponding database concepts. No matter how compelling and potentially valuable the vision may be, it is of little use until the legacy problem is solved. The problem of legacy systems migration is described, in the context of distributed computing, and is illustrated with lessons learned from actual case studies. The basic research directions and challenges are recast in the light of the reality of legacy systems. Recommendations for realizing the vision and meeting the challenges are given including the search for the elusive Killer Application.

## 1. World-Wide Computing

My professional goal is to contribute to making the world a better place by providing solutions to significant, real problems. As a computer science researcher this means that I want to produce the highest quality research and technology that is ultimately applicable to real problems so that the results are consistent with my beliefs. In this regard, I have very high hopes and expectations for the potential benefits of world-wide computing. The vision is that in world-wide computing the necessary computing resources (e.g., programs, information bases, information systems) can interact to cooperatively (or collaboratively) to solve problems effectively and efficiently and to do productive work. This will all happen transparently to the location or nature of the participating systems.

In this section, I describe a world-wide computing vision in terms of cooperation amongst information systems augmented by a telecommunications vision that provides communication on a scale previously unthinkable by computer scientists.

### 1.1. The Vision

The vision of distributed computing is compelling. It says that soon the dominant computing paradigm will involve large numbers of heterogeneous, intelligent agents distributed over large computer/communication networks. Agents may be humans, humans interacting with computers, humans working with computer support, and computer systems performing tasks without human intervention. Work will be conducted on the network in many forms. Work task definition will be centralized (e.g., a complex engineering task) and decentralized. Tasks will be executed by agents acting autonomously, cooperatively, or collaboratively, depending on the resources required to