

O. Lehrmann Madsen (Ed.)



ECOOP '92 European Conference on Object-Oriented Programming

Utrecht, The Netherlands, June 29-July 3, 1992
Proceedings

Springer-Verlag

Berlin Heidelberg New York
London Paris Tokyo
Hong Kong Barcelona
Budapest

BIBLIOTHEQUE DU CERIST

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 Editor

Ole Lehrmann Madsen
Computer Science Department, Aarhus University
Ny Munkegade, DK-8000 Aarhus C, Denmark

6239

CR Subject Classification (1991): D.2.2, D.3.2, H.2.4

ISBN 3-540-55668-0 Springer-Verlag Berlin Heidelberg New York
ISBN 0-387-55668-0 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

BIBLIOTHEQUE DU CERIST

Preface

This volume constitutes the proceedings of the sixth European Conference on Object-Oriented Programming (ECOOP), held in Utrecht (the Netherlands), June 29 - July 3, 1992. Previous ECOOP conferences were held in Paris (France), Oslo (Norway), Nottingham (England), Ottawa (Canada, jointly with OOPSLA), and Geneva (Switzerland). Next year's ECOOP is planned to take place in Kaiserslautern (Germany).

Since the "French initiative" to organize the first conference in Paris, ECOOP has been a very successful forum for discussing the state of art of object-orientation. ECOOP has been able to attract papers of a high scientific quality as well as high quality experience papers describing the pros and cons of using object-orientation in practice. This duality between theory and practice within object-orientation makes a good example of experimental computer science. The enormous commercial success of object-orientation may also be a problem for object-orientation. You cannot open a journal these days without reading about the great advantages of object-orientation. It is quite clear to me that many people are overselling object-orientation. The result of this may be that industry reacts negatively to object-orientation when it does not live up to all the given promises. We are sure that this year's ECOOP will continue the tradition of being a high quality conference and that it will contribute to a more balanced view on object-orientation.

This volume consists of 23 papers, including one invited paper and 22 papers selected by the programme committee from 124 submissions. The selection process was very difficult due to the large amount of good papers and many good papers could not be accepted due to space limitations. Each submitted paper has been reviewed by 3-4 people. The selection of papers was based only on the quality of the papers themselves. The quality of a conference is determined by the quality of the submitted papers and the quality of the reviewing. The members of the programme committee and the other referees have invested a large amount of work in the reviewing. The authors of submitted papers, the programme committee and the other referees deserve my sincere thanks.

The conference includes presentation of invited and selected papers, a panel, tutorials, workshops, demonstrations and an exhibition. In addition, the 25th anniversary of Simula 67, the first object-oriented language, will be celebrated. I would like to thank Mehmet Aksit, Pierre America, Gert Florijn, Paul Hendriks and Chris Laffra for the great job they have done in organizing this conference. I would also like to thank the participants for their contributions. Further thanks are due to the sponsoring organizations; and to Susanne Brøndberg, who was a great help in organizing the programme committee work. The success of ECOOP'92 is due to all these people and organizations.

May 1992

Ole Lehrmann Madsen
ECOOP'92 Programme Chair

BIBLIOTHEQUE DU CERIST

Organization

Conference Chair:	Pierre America (The Netherlands)
Programme Chair:	Ole Lehrmann Madsen (Denmark)
Organizing Chair:	Gert Florijn (The Netherlands)
Tutorials:	Mehmet Akşit (The Netherlands)
Exhibition:	Paul Hendriks (The Netherlands)
Demonstrations:	Chris Laffra (The Netherlands)

Sponsor

Software Engineering Research Centre - SERC, the Netherlands

Co-sponsoring Organizations

BCS (British Computer Society)	ICG
Datex	Koning en Hartman
DND (Norwegian Computer Society)	NGI (Dutch Computer Society)
Electric Engineering	Object Management Group
Georg Heeg Smalltalk-80-Systems	OCG (Austrian Computer Society)
Getronics Network Services	Philips
Getronics Service	SI (Swiss Computer Society)
IBM	Sun Microsystems

Programme Committee

Ole Lehrmann Madsen	Aarhus University, Denmark
Alan Borning	University of Washington, USA
Jean-Pierre Briot	Université Pierre et Marie Curie, France
Pierre Cointe	Université Pierre et Marie Curie, France
Elsbeth Cusack	British Telecom, U.K.
Nigel Derrett	Hewlett-Packard Laboratories, U.K.
Klaus Dittrich	Universität Zurich, Switzerland
Roger Duke	University of Queensland, Australia
Eric Juul	Københavns Universitet, Denmark
Boris Magnusson	Lund University, Sweden
Bertrand Meyer	Interactive Software Engineering, USA
Ron Morrison	University of St. Andrews, U.K.
Oscar Nierstrasz	Université de Genève, Switzerland
Walter Olthoff	German Research Center for AI, Germany
Barbara Pernici	Università di Udine, Italy
Jean-François Perrot	Université Pierre et Marie Curie, France
Anna-Kristin Präfrock	Siemens Nixdorf, Germany
Trygve Reenskaug	Taskon A/S, Norway
Markku Sakkinen	University of Jyväskylä, Finland
Dave Thomas	Object Technology International, Canada
Mario Tokoro	Sony CSL, Japan
Enn Tyugu	Institute of Cybernetics, Estonia
Rebecca Wirfs-Brock	Instantiations, USA
Akinori Yonezawa	University of Tokyo, Japan
Ed Yourdon	American Programmer, USA

List of Referees

Jay Almarode
Birger Andersen
Patrick Arnold
Cecily Bailes
Brian Barry
Barbara Beaudoin
Jean Bell
Dag Belsnes
Mats Bengtsson
Stephanie Bodoff
Isabelle Borne
Malik Bouabsa
Marc-Michael Brandis
Alfred Brown
Böhnlein
Craig Chambers
Bernadette Charron-Bost
Robert Cole
Derek Coleman
Richard Connor
Betsy Cordingley
Andreas Creppert
Michael Crowther
Quintin Cutts
Ole-Johan Dahl
Lars-Ove Dahlin
Laurent Dami
Christophe Dony
Denise Draper
Roger Duke
Nigel Edwards
Oren Etzioni
Ed Felten
Bjorn Freeman-Benson
David Freestone
Philippe Gautron
Stein Gjessing
Vera Goebel
Joseph Goguen
Patrick Goldsack
Thorsten Gorchs
Matthieu Goutet
Nicolas Graube
Peter Grogono
Hele-Mai Haav
Michel Habib
Fiona Hayes
Görel Hedin
Andrew Herbert
Alex Heung
Toshio Hirotsu
Ian Holland
Kohei Honda
Yutaka Ishikawa
Kent Johnson
Niels Christian Juul
Paul King
Graham Kirby
Jørgen Lindskov Knudsen
Karl-Heinz Koester
Dimitri Konstantas
Kai Koskimies
A. Kotz
Stein Krogdahl
Susanna Lambertini
Douglas Lea
Sylvie Lemarié
Loïc Lescaudron
Hank Levy
Peter Linington
Brendan Mahony
Jacques Malenfant
Roberto Maiocchi
Mihail Matskin
Satoshi Matsuoka
Arne Maus
Mike Milinkovich
Tatsuo Minohara
Gerhard Mueller
David Munro
Birger Møller-Pedersen
Jean-Marc Nerson
Thai Nguyen
Peter Nickolas
Else Nordhagen
David Notkin
John O'Connell
Tim O'Connor
Maria Orłowska
Olaf Owe
François Pachet
Jens Palsberg
Michael Papatomas
Jaan Penjam
Seppo Puuronen
Timothy Regan
Didier Rémy
Barbara Rieche
Renaud Rioboo
Gordon Rose
Steve Rudkin
Heikki Saikkonen
Hayssam Saleh
Anne Salvesen
Michael Schwartzbach
Graeme Smith
Akikazu Takeuchi
Magnus Taube
François-Xavier Testard-Vaillant
Dave Thomson
Claudio Trotta
Mark Utting
Tarmo Uustalu
Juha Vihavainen
Benjamin Volosh
Shigeru Watari
Andrew Watson
Phillip Williams
Shirley Williams
Alan Wills
Jeremy Wilson
Michael Wilson
Francis Wolinski
Phillip M. Yelland
Yauhiko Yokote

BIBLIOTHEQUE DU CERIST

Contents

On Unifying Relational and Object-Oriented Database Systems	1
<i>Invited paper, Won Kim (UniSQL, Austin Texas)</i>	
Import Is not Inheritance - Why We Need Both: Modules and Classes	19
<i>Clemens A. Szyperski (ETH Zürich)</i>	
Object-Oriented Multi-Methods in Cecil	33
<i>Craig Chambers (University of Washington, Seattle)</i>	
Aggregation in a Behaviour Oriented Object Model	57
<i>Thorsten Hartmann, Ralf Jungclauss, Gunter Saake (Technische Universität Braunschweig)</i>	
Reasoning and Refinement in Object-Oriented Specification Languages	78
<i>K. Lano (Oxford University), H. Haughton (Lloyds Register of Shipping, Croydon)</i>	
Combining Object-Oriented and Logic Paradigms: A Modal Logic Programming Approach	98
<i>Tarmo Uustalu (Estonian Academy of Sciences, Tallinn)</i>	
An Incremental Class Reorganization Approach	114
<i>Eduardo Casais (Forschungszentrum Informatik, Karlsruhe)</i>	
System Design by Composing Structures of Interacting Objects	133
<i>Egil P. Andersen and Trygve Reenskaug (University of Oslo)</i>	
Unifying the Design and Implementation of User Interfaces Through the Object Paradigm	153
<i>Véronique Normand (Bull-IMAG Systèmes, Gières), Joëlle Coutaz (Laboratoire de Génie Informatique, Grenoble)</i>	
Nesting Actions Through Asynchronous Message Passing: The ACS Protocol .	170
<i>Rachid Guerraoui, Riccardo Capobianchi, Agnes Lanusse, Pierre Roux (CE Saclay DEIN/SIR, Gif sur Yvette)</i>	
Inheritance of Synchronization Constraints in Concurrent Object-Oriented Programming Languages	185
<i>Svend Frølund (University of Illinois, Urbana)</i>	
EPEE: an Eiffel Environment to Program Distributed Memory Parallel Computers	197
<i>J.M. Jézéquel (IRISA, Rennes)</i>	

Using Object-Oriented Programming Techniques for Implementing ISDN Supplementary Services	213
<i>Bent Gabelgaard (Dansk MobilTelefon, Aalborg)</i>	
An Object Model for Engineering Design	233
<i>G.T. Nguyen, D. Rieu, J. Escamilla (IRIMAG, Grenoble)</i>	
An Object-Oriented Class Library for Scalable Parallel Heuristic Search	252
<i>Wing-cheong Lau (University of Texas, Austin), Vineet Singh (Hewlett-Packard Labs, Palo Alto)</i>	
Integrating Constraints with an Object-Oriented Language	268
<i>Bjorn N. Freeman-Benson (University of Victoria), Alan Borning (University of Washington)</i>	
Specifying Reusable Components Using Contracts	287
<i>Ian M. Holland (Northeastern University, Boston)</i>	
ACTS: A Type System for Object-Oriented Programming Based on Abstract and Concrete Classes	309
<i>Mahesh Dodani, Chung-Shin Tsai (University of Iowa)</i>	
Making Type Inference Practical	329
<i>Nicholas Ozhøj, Jens Palsberg, Michael I. Schwartzbach (Aarhus University)</i>	
A Reflective Model of Inheritance	350
<i>Suresh Jagannathan (NEC Research Institute, Princeton), Gul Agha (University of Illinois, Urbana)</i>	
An Object-Oriented Language-Database Integration Model: The Composition-Filters Approach	372
<i>Mehmet Akşit, Lodewijk Bergmans, Sinan Vurul (University of Twente)</i>	
Supporting Physical Independence in an Object Database Server	396
<i>Nicola Aloia (CNUCE-CNR, Pisa), Svetlana Barneva, Fausto Rabitti (IEI-CNR, Pisa)</i>	
Developing a Class Hierarchy for Object-Oriented Transaction Processing	413
<i>Daniel L. McCue (University of Newcastle upon Tyne)</i>	

On Unifying Relational and Object-Oriented Database Systems

Won Kim

UniSQL, Inc.
9390 Research Blvd.
Austin, Texas 78759

Abstract. The past several years have been the gestation period for a new generation of database technology. There has been a flurry of activities to develop and experiment with database systems that support an object-oriented data model or that extend the relational data model with some object-oriented facilities. These activities have been fueled by the emergence of a broad spectrum of database applications which relational database systems cannot support and the increasing need to achieve another productivity leap in application development. As a result of these efforts, there is now a sufficient body of knowledge for the development of a commercially viable next-generation database system. Such a system should support a unified relational and object-oriented data model; that is, a full object-oriented data model in a way that is completely compatible with the relational model. This article examines motivations for unifying the relational and object-oriented data models in a single database system, and outlines design and implementation issues that must be addressed in building such a system.

1 Introduction

During the 1980s, relational database systems [11, 12, 13, 33] dominated the database market for business data processing applications. The relational database language SQL [3] became an industry standard. The theory, implementation, and use of database systems became a major discipline of computer science. The simplicity of the relational model of data and the dynamic management of a database have been accepted as vehicles for significant productivity enhancement in application development.

However, even as the acceptance of relational database systems spread, their limitations were exposed by the emergence of various classes of new applications. These applications include multimedia systems, statistical and scientific modeling and analysis systems, geographic information systems, engineering and design systems, knowledge-based systems, and so on. The limitations of relational database systems that these applications exposed fall into two categories: data model and computational model. The data-modeling facilities that relational systems lack include those for specifying, querying and updating

BIBLIOTHEQUE DU CERIST

complex nested entities (such as design and engineering objects, and compound documents); arbitrary user-defined data types; frequently useful relationships, such as generalization and aggregation relationships; temporal evolution of data (i.e., temporal dimension of data, and versioning of data); and so on. The computation-modeling facilities that relational systems lack include the management of memory-resident objects for extensive pointer-chasing applications (e.g., simulation of a computer-aided design); long-duration, cooperative transactions; and so on.

The need to reduce the cost of developing and operating these applications pointed to the need for a fundamental advancement in database technology, that is, a paradigm shift, rather than incremental extensions of the capabilities of existing data management systems. The basis of this fundamental advancement in database technology is the object-oriented paradigm developed in object-oriented programming languages. There are two major reasons the object-oriented paradigm is a sound basis for a new generation of database technology.

One is that the object-oriented paradigm [16, 29, 32] can be the basis for a data model which subsumes the relational (and pre-relational) data models. Solutions to most of the data-modeling-related difficulties of relational database systems are inherent in an object-oriented data model. Relational systems are designed to manage only limited types of data, such as integer, floating-point number, string, Boolean, date, time, and money. In other words, they are not designed to manage arbitrary user-defined data types. On the other hand, a central tenet of an object-oriented data model is the uniform treatment of arbitrary data types and the facility to add new data types. Further, an object-oriented data model allows the representation of not only data, and relationships and constraints on the data, as the relational data model does; but also the encapsulation of data and programs that operate on the data, and provides a uniform framework for the treatment of arbitrary user-defined data types.

Another reason is that the object-oriented paradigm, through the notions of encapsulation and inheritance (rusec), is fundamentally designed to reduce the difficulty of developing and evolving complex software systems or designs. This is precisely the goal that has driven the data management technology from file systems to relational database systems during the past three decades. An object-oriented data model thus inherently satisfies the objective of reducing the difficulty of design and evolving very large and complex databases. The notions of encapsulation and inheritance are a key to the search for further productivity enhancement in database application development.

The promises of a database technology based on an object-oriented data model are clear. These promises fueled, during the past several years, a rush to develop a first wave of object-oriented database systems [30, 2, 8, 15, 14, 20]. About a half dozen systems are currently commercially available; some are persistent storage systems for C++ applications, some provide low-level support for engineering and design applications; and some offer a relatively richer set of database features [25].

If the object-oriented technology can truly deliver the fundamental advance in database technology to the overall database market, that is, transition the database technology past the

current relational technology, it needs to be unified with the relational technology. The relational paradigm should be extended with a set of fundamental object-oriented concepts found in most object-oriented programming languages and application systems. The database language that embodies the united modeling paradigm should be a minimal extension to the SQL relational database language. The database language should then be embedded in a wide variety of host programming languages to bring object-oriented modeling facilities to programmers of these languages.

SQL-compatibility is a key to the adoption of new database systems that incorporate object-oriented modeling facilities. It will minimize the time necessary for the current relational users to learn the new technology, and facilitate the migration of relational applications to the new systems. However, SQL-compatibility will not solve the mismatch that has existed until now between a database language and the host programming language in which the database language is embedded. Ideally, database application programmers should use a single database programming language for both general-purpose programming and database management. Once object-oriented programming languages become solidly established as database application programming languages, database systems that "seamlessly" support the languages will certainly be useful.

Some vendors currently offer unified database programming languages by augmenting object-oriented programming languages, notably, C++ and Smalltalk, with persistent storage, transaction management, and limited query facilities. However, important challenges remain. Persistent storage and database management facilities (especially the query facilities) must be provided without introducing obtrusive syntax or semantics that is different from that of the programming language being augmented. Database management facilities and persistent storage should be provided to objects of all data types allowed in the programming language. Further, database management facilities should be provided to not only persistent objects, but also nonpersistent objects. Finally, the extensions to one programming language should be applicable with minimal changes to a variety of other programming languages. It is important to note that, regardless of whether the database language is an extension of SQL or some object-oriented programming language, most of the considerations outlined in this paper remain valid. Ultimately, a database system must address the impact of object-oriented modeling concepts on the database architecture.

A formidable array of technical challenges must be overcome before a next-generation database system can be built that incorporates object-oriented modeling concepts into relational database systems. The purpose of this paper is to outline these challenges and how they may be met. First, the relational and object-oriented data models must be combined into a single coherent unified model. Second, a database language must be designed to allow the specification of data and relationships among them, and the querying and updating of the data. Third, the database system must be built to fully support all the facilities the database language allows. Fourth, the system must deliver high performance; in particular, it must deliver performance comparable to relational database systems for all operations that can be performed using relational database systems. The strength of an object-oriented data model is also its weakness; the richness of an object-oriented data model that makes it possible to

model complex objects and their relationships in nonbusiness data-processing applications necessarily introduces added complexities with which the users must cope. The richness of an object-oriented data model also implies added difficulties in implementing an object-oriented database system. The designers of the unified database language and database system must carefully revisit the philosophy and design rationale behind relational database systems, and address a host of technical problems in database architecture.

2 Design Issues

A database system is in essence a software that implements all the functions supported in a database language. A database language in turn is an embodiment of a data model. For example, the relational database language SQL stipulates how the users of a database system should create tables; specify data types and integrity constraints on the tables and columns within tables; populate the tables with rows; query, update, and delete the contents of the database; and so forth. A data model is thus the foundation of any database system.

Further, a database language consists of three sublanguages: data-definition, query and data manipulation, and data control language. A data-definition language allows the definition of a database. A query and data manipulation language allows the database to be populated, and the database contents to be queried, updated, and deleted. A data control language allows the specification of database sharing and administration.

In this section, I will show first that a unified relational and object-oriented data model is in essence an object-oriented data model obtained by extending the relational model with key concepts found in object-oriented programming. Next, I will outline in terms of the three database languages architectural issues in building a unified database system.

2.1 Data Model

The primary advantage of the relational model of data is its simplicity; however, simplicity is also a major disadvantage of the relational model. A relational database consists of a set of relations (tables), and a relation in turn consists of rows and columns. An entry in a table may have a single value, and the value may belong to a set of system-defined data types (e.g., integer, string, float, date, time, money). The user may impose further restrictions, called integrity constraints, on these values (e.g., the integer value of an employee age may be restricted to between 18 and 65).

This simple data model is really a subset of an object-oriented data model; that is, the relational model can be generalized to arrive at an object-oriented data model. First, let us regard a table as a data type, and allow each entry of a table to be a single value belonging to any arbitrary user-defined table, rather than just the system-defined data types. The first CREATE TABLE statement in Figure 1 shows the specification of an Employee table under the relational model. The values of the Hobby and WorksFor columns are restricted to character strings. The second CREATE TABLE in Figure 1 reflects data-type generalization for the columns of a table. The value for the Hobby column no longer needs to be restricted to a character string; it may now be a row of a user-defined table Activity.