

Paolo Atzeni (Ed.)

Cco 1-701

LOGIDATA+: Deductive Databases with Complex Objects

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
D-76131 Karlsruhe, FRG

Juris Hartmanis
Cornell University
Department of Computer Science
4130 Upson Hall
Ithaca, NY 14853, USA

Volume Editor

Paolo Atzeni
Dipartimento di Informatica e Sistemistica
Terza Università di Roma and Università di Roma "La Sapienza"
Via Salaria, 113, I-00198 Rome, Italy

CR Subject Classification (1991): H.2.1, H.2.3

6303

ISBN 3-540-56974-X Springer-Verlag Berlin Heidelberg New York
ISBN 0-387-56974-X 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
45/3140-543210 - Printed on acid-free paper

Preface

This book presents a collection of coordinated scientific papers describing the work conducted and the results achieved within the LOGIDATA+ project, a research action funded by CNR (Consiglio Nazionale delle Ricerche – the Italian National Research Council), within "Progetto Finalizzato Sistemi Informatici e Calcolo Parallelo."

The aim of the LOGIDATA+ project is the definition of advanced database systems that significantly extend the functionalities of the current systems, with specific reference to the application areas for which relational systems are not considered satisfactory. These new systems will allow the definition of data with complex structures, the representation of semantic relationships between objects, and the use of powerful query and update languages. They will be based on a combination of techniques originating from relational databases and logic programming, with contributions from object-oriented programming. Attention will be devoted to the representation of isa-hierarchies and to taxonomic reasoning. The specific goal of the project is the definition, design, and prototype implementation of a database management system with complex structures and a class hierarchy, to be accessed through a rule-based language.

The project is now at the conclusion of the first phase, with significant research work done on the definition of the features of the systems and their theoretical foundations. The second phase will be mainly concerned with the implementation of prototypes. This book presents a homogeneous, integrated view of the scientific results of the project, with respect to all the features of the system.

The LOGIDATA+ group involves teams from the following institutions:

- CNR, Centro per l'Interazione Operatore-Calcolatore, Bologna
- CNR, Istituto di Analisi dei Sistemi ed Informatica, Roma
- CRAI, Rende (Cosenza)
- Politecnico di Milano, Dipartimento di Elettronica e Informazione
- Sintesi S.r.l., Modena
- Systems & Management S.p.A., Torino
- Università dell'Aquila, Dipartimento di Matematica Pura e Applicata
- Università della Calabria, Cosenza, Dipartimento di Elettronica, Informatica e Sistemistica
- Università di Firenze, Dipartimento di Sistemi e Informatica
- Università di Modena, Dipartimento di Matematica
- Università di Roma La Sapienza, Dipartimento di Informatica e Sistemistica

The book is organized in four parts. Part I contains two papers, the first giving an overview of the LOGIDATA+ project and the second presenting a general discussion on the combination of deductive and object-oriented features in the database field. Part II is concerned with the description of the LOGIDATA+ model and language: the first and second papers present the data model and the rule-based language, respectively; the third paper illustrates the concepts by means of the detailed description of an application. Parts III and IV report on research results about a number of issues that can lead to significant extensions of the LOGIDATA+ system. Each of them concentrates on some features of the model and language. Specifically, Part III contains results on problems related to structural issues (updates over classes, taxonomic reasoning, and integrity constraints in object-oriented databases) and Part IV deals with deductive issues, essentially extensions of logic programming. Finally, Part V presents the experimental results of the project, with the existing prototypes.

I would like to thank all those who have made this book possible. Bruno Padini, Director of "Progetto Finalizzato Sistemi Informatici e Calcolo Parallelo," and Domenico Sacca, Coordinator of "Sottoprogetto 5: Sistemi evoluti per basi di dati," have encouraged the LOGIDATA+ action since its inception. The authors of all the papers deserve my gratitude for their timely cooperation in the preparation of the volume and for their help in the revision process. My graduate students Giovanni Barone, Luca Cabibbo, and Gianni Mecca helped in the reviewing process and in the organization of the material. Alfred Hofmann at Springer-Verlag was very efficient in the management of our project.

May 1993

Paolo Atzeni

Table of Contents

Part I: The Framework

LOGIDATA+: Overview	
<i>Paolo Atzeni</i>	1
Bridging Objects with Logical Rules: Towards Object-Oriented Deductive Databases	
<i>Stefano Ceri and Letizia Tanca</i>	7

Part II: The Model and the Language

The LOGIDATA+ Model	
<i>Paolo Atzeni, Filippo Carace, Stefano Ceri, and Letizia Tanca</i>	20
The LOGIDATA+ Language and Semantics	
<i>Paolo Atzeni, Luca Cabibbo, Giansalvatore Mecca, and Letizia Tanca</i>	30
Travel Agency: A LOGIDATA+ Application	
<i>Luca Cabibbo and Giansalvatore Mecca</i>	42

Part III: Complex Objects

Management of Extended Update Operations	
<i>Luigi Palopoli and Riccardo Torlone</i>	60
Taxonomic Reasoning in LOGIDATA+	
<i>Domenico Beneventano, Sonia Bergamaschi, Claudio Sartori, Alessandro Artale, Francesca Cesarini, and Giovanni Soda</i>	79
Introducing Taxonomic Reasoning in LOGIDATA+	
<i>Alessandro Artale, Francesca Cesarini, and Giovanni Soda</i>	85
Taxonomic Reasoning with Cycles in LOGIDATA+	
<i>Domenico Beneventano, Sonia Bergamaschi, and Claudio Sartori</i>	105
Modeling Semantic Integrity Constraints in Object-Oriented Database Schemas	
<i>Anna Formica and Michele Missikoff</i>	129

Part IV: Deductive Databases

Evaluation of Negative Logic Programs

Sergio Greco, Massimo Romeo, and Domenico Saccà 148

Effective Implementation of Negation in Database Logic

Query Languages

Nicola Leone, Massimo Romeo, Pasquale Rullo, and Domenico Saccà 159

Modules in Logic Programming:

A Framework for Knowledge Management

Annalina Fabrizio, Maurizio Capaccioli, and Sandra Valeri 176**Part V: Prototypes**

LOA: The LOGIDATA+ Object Algebra

Umberto Nanni, Silvio Salza, and Mario Terranova 195

The LOGIDATA+ Prototype System

Umberto Nanni, Silvio Salza, and Mario Terranova 211MOOD*: An Architecture for Object-Oriented Access
to a Relational Database*Marco Lugli, Luca Nini, and Stefano Ceri* 233

Prototypes in the LOGIDATA+ Project

*Alessandro Artale, Jean-Pierre Ballerini, Sonia Bergamaschi,
Filippo Cacace, Stefano Ceri, Francesca Cesarini, Anna Formica,
Herman Lam, Sergio Greco, Giorgio Marrella, Michele Missikoff,
Luigi Palopoli, Luigi Pichetti, Domenico Saccà, Silvio Salza,
Claudio Sartori, Giovanni Soda, Letizia Tanca, and Marco Toiati* 252

LOGIDATA+: Overview*

Paolo Atzeni

Dipartimento di Informatica e Sistemistica, Università di Roma "La Sapienza,"
Via Salaria 113, 00198 Roma, Italy

Abstract. LOGIDATA+ is part of the subproject "Advanced Database Systems" of the project "Information Systems and Parallel Computation", of the Italian research council. The aim of LOGIDATA+ is the definition of advanced database systems, which significantly extend the functionalities of the current systems, with specific reference to the application areas for which relational systems are not considered satisfactory. These new systems will allow the definition of data with complex structures, the representation of semantic relationships between objects, and the use of powerful query and update languages. They will be based on the combination of techniques originating from relational databases and logic programming, with contributions from object-oriented programming. Attention will be devoted to the representation of isa-hierarchies and to taxonomic reasoning. The specific goal of LOGIDATA+ is the definition, design, and prototypal implementation of a database management system with a complex structure and a class hierarchy, to be accessed through a rule-based language. Extensions of the core are also considered, in various directions: management of various forms of negation, both in the body and in the head of the rules; management of updates, especially with respect to intensionally defined predicates or to rules; modularization; constraints; taxonomic reasoning.

1 Introduction

LOGIDATA+ is a research action aimed at the definition of advanced database systems, which significantly extend the functionalities of the current systems, with specific reference to the application areas for which relational systems are not considered satisfactory. These new systems will allow the definition of data with complex structures, the representation of semantic relationships between objects, and the use of powerful query and update languages. They will be based on the combination of techniques originating from relational databases and logic programming, with contributions from object-oriented programming. Attention will be devoted to the representation of isa-hierarchies and to taxonomic reasoning.

The specific goal of LOGIDATA+ is the definition, design, and prototypal implementation of a database management system with a complex structure and a class hierarchy, to be accessed through a rule-based language.

* Work supported by CNR, "Progetto Finalizzato Sistemi Informatici e Calcolo Parallelo." The author is now with Terza Università di Roma.

LOGIDATA+ is part of a larger national project, on Information Systems and Parallel Computation (officially, "Sistemi Informatici e Calcolo Parallelo," abbreviated as SICP). More specifically, it is part of a subproject on Advanced Database Systems of the SICP Project. This paper gives a brief overview of the LOGIDATA+ action and of the framework in which it operates. First we describe the SICP Project (Section 2) and its subproject on databases (Section 3). Then we discuss the motivation and the goals of the LOGIDATA+ action, and finally we describe its organization and the way its results are described in the subsequent papers of this book (Section 4).

2 The Project "Sistemi Informatici e Calcolo Parallelo"

In the Italian system, a *Progetto Finalizzato del CNR* is "a set of research activities aimed at reaching goals that have a significant social and economic interest at the national level, with the involvement of all the components of the scientific system of the country." A *Progetto Finalizzato* is funded by the national government and administered by *Consiglio Nazionale delle Ricerche* (CNR, the Italian research council).

The Project *Sistemi Informatici e Calcolo Parallelo* (SICP) was approved by the government in 1987 together with other nine projects. Their general goal is to support the national industry in the development and acquisition of new technology, in order to increase its competitiveness on the international market and to reduce the need for technology import. They also involve the development of prototypes, the study and evaluation of existing system and the definition of proposals for the governmental activity in various fields.

The total budget of the SICP Project (throughout the five-year length) is 63 billion Lire (about 40 million US dollar, at the current exchange rate). The project is organized in seven subprojects, over various areas of computer science and engineering, plus a support initiative for the coordination of laboratories:

- SP1 Scientific computation for large systems
- SP2 Dedicated processors
- SP3 Parallel architectures
- SP4 New generation languages
- SP5 Advanced database systems
- SP6 Methods and tools for system design
- SP7 Support systems for intellectual work
- SP8 Support initiative

The Director of the project is Bruno Fadini, Università di Napoli "Federico II." The specific goals of each subproject include both research results that contribute to the advancement of knowledge in the field and more practical results, such as methods, systems, and services. The activities have been divided in two phase: the first three-year period had mainly research goals with some experimental work whereas the final two year (just started) will mainly be devoted to the finalization of the results, with significant prototypes.