DATABASE TRANSACTION MODELS

For Advanced Applications

Edited

by

AHMED K. ELMAGARMID

Database Transaction Models for Advanced Applications

Edited by

Ahmed K. Elmagarmid

MORGAN KAUFMANN PUBLISHERS
San Mateo, California

Editor: *Bruce M. Spatz*
Production Manager: *Yonie Overton*
Cover Mechanical Artist: *Patty King*
Indexing/Electronic Output: *Superscript Typography*

Morgan Kaufmann Publishers, Inc.
Editorial Office:
2929 Campus Drive, Suite 260
San Mateo, CA 94403

95 94 93 92 5 4 3 2 1

Library of Congress Cataloging in Publication Data is available for this book.

ISBN 1-55860-214-3
ISSN 1046-1698

Database Transaction Models for Advanced Applications

THE MORGAN KAUFMANN SERIES IN
DATA MANAGEMENT SYSTEMS

Series Editor, Jim Gray

DATABASE TRANSACTION MODELS FOR ADVANCED APPLICATIONS
Edited by Ahmed K. Elmagarmid (Purdue University)

A GUIDE TO DEVELOPING CLIENT/SERVER SQL APPLICATIONS
Setrag Khoshafian (Portfolio Technologies, Inc.),
Arvola Chan (Versant Object Technology),
Anna Wong (CLaM Associates), and
Harry K. T. Wong (Nomadic Systems)

THE BENCHMARK HANDBOOK FOR DATABASE AND TRANSACTION PROCESSING SYSTEMS
Edited by Jim Gray (Digital Equipment Corporation)

CAMELOT AND AVALON: A DISTRIBUTED TRANSACTION FACILITY
Edited by Jeffrey L. Eppinger (Transarc Corporation),
Lily B. Mummert (Carnegie Mellon University),
and Alfred Z. Spector (Transarc Corporation)

DATABASE MODELING AND DESIGN: THE ENTITY-RELATIONSHIP APPROACH
Toby J. Teorey (University of Michigan)

READINGS IN OBJECT-ORIENTED DATABASE SYSTEMS
Edited by Stanley B. Zdonik (Brown University) and
David Maier (Oregon Graduate Center)

READINGS IN DATABASE SYSTEMS
Edited by Michael Stonebraker (University of California at Berkeley)

DEDUCTIVE DATABASES AND LOGIC PROGRAMMING
Jack Minker (University of Maryland)

To beautiful little
Najd Mohammed Elmagarmid,
April 7, 1988–December 15, 1990

Foreword

The transaction concept has emerged as the key structuring technique for distributed data and distributed computations. Originally developed and applied to database applications, the transaction model is now being used in new application areas ranging from process control to cooperative work. Not surprisingly, these more sophisticated applications require a refined and generalized transaction model. The concept must be made recursive, it must deal with concurrency within a transaction, it must relax the strict isolation among transactions, and it must deal more gracefully with failures.

This book collects, for the first time, the many generalizations of the basic transaction model. Taken together it represents the latest and best thinking on how to structure distributed and heterogeneous computations. The authors represent the spectrum from research universities to industry. The applications they discuss encompass planning, design, engineering, manufacturing, and commerce. Their approaches range from theoretical models to practical tool kits. In sum, this book covers the entire spectrum of thought on the topic of generalized transaction models and transaction processing.

This book stands as a milestone, collecting the many intellectual threads developed over the decade of the 1980s. Researchers are now working to merge these many approaches into a new unified transaction concept able to handle all applications. This new emerging concept will likely be one of the major contributions to Computer Science in the coming decade.

Jim Gray
Digital Equipment Corporation

Preface

What constitutes a transaction model? This innocent question from an uninformed colleague of mine stopped me, for the first time in ten years, to really think about this issue. It is easy to understand that a data model describes the structure of the database, i.e., data types and the constraints related to the data. But what components one should find in a transaction model? After mumbling something about how to specify the correctness of the database, and the type and structure of the logical operations to manipulate it, I quickly realized that giving a satisfactory answer is not as easy I thought. Several more hours of pondering upon the topic gave more refined, but equally unsatisfactory answers - especially in the light of the new developments in the transaction processing research. All my attempts to give a definition centered not only on the description of properties of the notion of a transaction, but also on the external mechanisms that guarantee such properties. This is exactly what the different chapters of this book are all about.

Active research on database transaction processing dates back to the early '70s, and many theoretical and pragmatic issues of the "traditional transactions" have been explored. The key properties such as atomicity of transaction execution, isolation of intermediate results and a universal correctness criteria (serializability) have been discovered and studied intensively. However, as pointed out by many early researchers, much of this discussion was intended to be relevant for business oriented database applications, where transactions were assumed to be simple, short-lived and non-hierarchical in structure. Attempts to blindly apply this work to other contexts has almost in every case that I know of, produced impractical and sometimes even disastrous results. Indeed, one-by-one each one of the fundamental "ACID" properties has been challenged in various environments, in spite of the appeal of the theoretical elegance of such concepts as serializability or atomic commitment.

So what are the new issues that need to be addressed by these new

transaction models? This many faceted question is best answered by reading the various chapter in this book; my attempt here is only to illustrate how a single modification in the assumptions gives rise to several major questions to be answered.

Let's look at one simple (realistic) modification that the transactions are long-lived, their duration is not seconds but minutes, hours or weeks. This type of transactions are especially natural in Computer Aided Design (CAD) or cooperative environments (e.g., in Groupware). The most common commercially used concurrency control method, blocking the data via locks, becomes totally impractical as one would exclude all the access to data for hours. Actually this single modification severely challenges the principles of guaranteeing isolation and atomicity of the traditional transactions, and introduces naturally the notion of compensation used in S-transactions, Sagas, Open Nested transactions and transactions on Active Objects. Secondly, as such environments are decentralized in nature, it is natural to extend the traditional flat transaction control structure to a hierarchical one, i.e., introduce the notion of transaction nesting with subtransactions executing at various sites. Finally, for such long-living transactions it is also typical to allow several alternative execution threads which leads to incorporating multiple acceptable final states into the correctness criteria (e.g., Flex transactions).

The different chapters of this book provide an excellent snapshot of the various questions raised by the new application areas, where transaction processing is needed, but where the underlying assumptions make it impossible to apply the traditional models and algorithms. The book acts both as a reference for the state of the art in advanced transaction models, as well as a good starting point for anybody interested in addressing the important open questions in the area. Since the chapters are written by different authors, the book has unique value in collecting to a single volume various alternative views on the topics, fact that made the book especially appealing to me personally.

Henry Tirri
University of Helsinki

Acknowledgements

This book would not have come about without the contributions and dedication of the many authors involved. This book is an acknowledgement of their many and valuable contributions to the field.

The contributions have endured a roller coaster that lasted for over two years. What was promised to be a fast and painless task took five times longer than promised and went through far too many revisions. We all are indebted to the over 100 reviewers who did the hard job. I was pleased at the level of professionalism and desire for the highest levels of perfection exhibited by the contributors and reviewers.

This book would not have been possible without the help of many of the InterBase team at Purdue. A special mention goes to Sue Barton, Omran Bukhres, Jiansan Chen, Yungho Leu, James Mullen, and Aidong Zhang.

A great amount of input was provided by so many of my good friends. I will not list them all, but a special mention is given to Amit Sheth and M. Tamer Özsu. Special thanks are due to Henry Tirri for writing a preface to the book and to Jim Gray, the series editor, for writing the Foreword to the book.

Special thanks go to the staff at Morgan Kaufmann. Their dedication to this project has been unyielding. Special thanks go to my Production Manager Yonie Overton. She has put up with a tight and often unpredictable schedule. No amount of thanks or appreciation is enough for the force behind the project, Bruce Spatz, the Senior Editor for this book. Indeed he is solely responsible for the inception of the idea of writing the book on transaction models and receives all the credit for seeing the project through thick and thin. He has been gracious, friendly and accommodating at times when I may have been unreasonable. Thanks are due also to Elizabeth Essex, Associate Product Manager.

The push for writing a book started at a meeting with Mike Morgan in

my office at Purdue in the Autumn of 1988. A lot has happened in the field since that time, and we hope we have captured some of it in these pages.

This book does not particularly cover heterogeneous databases, as that is the topic of an upcoming book by myself, Amit Sheth and Marek Rusinkiewicz.

This book would not have been possible without the financial support I receive from the Purdue Research Foundation, The Indiana Business Modernization and Technology, The National Science Foundation, The Software Engineering Research Center, Bell Northern Research, and NASA.

Ahmed K. Elmagarmid
Associate Professor and Executive Director
Indiana Center for Database Systems
Department of Computer Sciences
Purdue University

Contents