IMPRECISE AND APPROXIMATE COMPUTATION

THE KLUWER INTERNATIONAL SERIES IN ENGINEERING AND COMPUTER SCIENCE

REAL-TIME SYSTEMS

Consulting Editor

John A. Stankovic

FOUNDATIONS OF DEPENDABLE COMPUTING: ***System Implementation,*** edited by Gary M. Koob and Clifford G. Lau, ISBN: 0-7923-9486-0

FOUNDATIONS OF DEPENDABLE COMPUTING: ***Paradigms for Dependable Applications,*** edited by Gary M. Koob and Clifford G. Lau, ISBN: 0-7923-9485-2

FOUNDATIONS OF DEPENDABLE COMPUTING: ***Models and Frameworks for Dependable Systems,*** edited by Gary M. Koob and Clifford G. Lau, ISBN: 0-7923-9484-4

THE TESTABILITY OF DISTRIBUTED REAL-TIME SYSTEMS, Werner Schütz; ISBN: 0-7923-9386-4

A PRACTITIONER'S HANDBOOK FOR REAL-TIME ANALYSIS: ***Guide to Rate Monotonic Analysis for Real-Time Systems,*** Carnegie Mellon University (Mark Klein, Thomas Ralya, Bill Pollak, Ray Obenza, Michale González Harbour); ISBN: 0-7923-9361-9

FORMAL TECHNIQUES IN REAL-TIME FAULT-TOLERANT SYSTEMS, J. Vytopil; ISBN: 0-7923-9332-5

SYNCHRONOUS PROGRAMMING OF REACTIVE SYSTEMS, N. Halbwachs; ISBN: 0-7923-9311-2

REAL-TIME SYSTEMS ENGINEERING AND APPLICATIONS, M. Schiebe, S. Pferrer; ISBN: 0-7923-9196-9

SYNCHRONIZATION IN REAL-TIME SYSTEMS: ***A Priority Inheritance Approach,*** R. Rajkumar; ISBN: 0-7923-9211-6

CONSTRUCTING PREDICTABLE REAL TIME SYSTEMS, W. A. Halang, A. D. Stoyenko; ISBN: 0-7923-9202-7

FOUNDATIONS OF REAL-TIME COMPUTING: ***Formal Specifications and Methods,*** A. M. van Tilborg, G. M. Koob; ISBN: 0-7923-9167-5

FOUNDATIONS OF REAL-TIME COMPUTING: ***Scheduling and Resource Management,*** A. M. van Tilborg, G. M. Koob; ISBN: 0-7923-9166-7

REAL-TIME UNIX SYSTEMS: ***Design and Application Guide,*** B. Furht, D. Grostick, D. Gluch, G. Rabbat, J. Parker, M. McRoberts, ISBN: 0-7923-9099-7

IMPRECISE AND APPROXIMATE COMPUTATION

edited
by

Swaminathan Natarajan
Xerox Corporation

KLUWER ACADEMIC PUBLISHERS
Boston / Dordrecht / London

Distributors for North America:
Kluwer Academic Publishers
101 Philip Drive
Assinippi Park
Norwell, Massachusetts 02061 USA

Distributors for all other countries:
Kluwer Academic Publishers Group
Distribution Centre
Post Office Box 322
3300 AH Dordrecht, THE NETHERLANDS

Library of Congress Cataloging-in-Publication Data

A C.I.P. Catalogue record for this book is available
from the Library of Congress.

Printed on acid-free paper.

Printed in the United States of America

CONTENTS

LIST OF FIGURES

Chapter 1

Chapter 2

LIST OF TABLES

CONTRIBUTORS

Paul D. Alexander
Department of Electrical &
Electronic Engineering,
University of Adelaide, Australia

N. C. Audsley
Real-Time Systems Research Group,
Department of Computer Science,
University of York,
York, United kingdom

Alan Burns
Real-Time Systems Research Group,
Department of Computer Science,
University of York,
York, United kingdom

Robert I. Davis
Real-Time Systems Research Group,
Department of Computer Science,
University of York,
York, United kingdom

Alan Garvey
Department of Computer and Information Science,
University of Massachusetts,
Amherst, Massachusetts 01003

Victor Lesser
Department of Computer and Information Science,
University of Massachusetts,
Amherst, Massachusetts 01003

Joseph Y-T. Leung
Department of Computer Science
and Engineering,
University of Nebraska-Lincoln,
Lincoln, Nebraska 68588-0115

Cheng Chew Lim,
Department of Electrical &
Electronic Engineering,
University of Adelaide, Australia

Kwei-Jay Lin
Department of ECE,
University of California,
Irvine, California 92715

Jane W. S. Liu
Department of Computer Science,
University of Illinois,
Urbana, Illinois 61801

Swaminathan Natarajan
Xerox Corporation,
1350 Jefferson Road,
Henrietta, New York 14623

Stuart Russell
Department of EECS,
Computer Science Division,
University of California,
Berkeley, California 94720

Alvaro del Val
Robotics Lab,
Computer Science Department,
Stanford University,
Stanford, California 94305

Susan V. Vrbsky
Department of Computer Science,
University of Alabama,
Tuscaloosa, Alabama 35487

A. J. Wellings
Real-Time Systems Research Group,
Department of Computer Science,
University of York,
York, United kingdom

John Yen
Department of Computer Science,
Texas A&M University,
College Station, Texas 77843-3112

Albert C. Yu
Systems Division,
RE/R8/N541,
Hughes Aircraft Company,
El Segundo, California 90245

Wei Zhao
Department of Computer Science,
Texas A&M University,
College Station, Texas 77843-3112

Shlomo Zilberstein
Department of Computer Science,
University of Massachusetts,
Amherst, Massachusetts 01003

PREFACE

Real-time systems are now used in a wide variety of applications, including space systems, defense systems, process control and signal processing. Conventionally, real-time systems were configured at design time to perform a given set of tasks, and could not adapt easily to dynamic situations. Recently, the concept of imprecise and approximate computation has emerged as a promising approach to providing scheduling flexibility and enhancing dependability in dynamic real-time systems. This volume is a collection of papers which chart the current state of the art in imprecise and approximate computation.

In real-time systems, the focus is on ensuring that a set of tasks each complete within their deadlines. Meeting deadlines is critical in a real-time system, because these systems typically interact with the external environment, and a missed deadline may lead to a system failure which causes a catastrophic situation in the environment. For example, if the obstacle detection task in a moving robot misses its deadline, the robot may crash into a wall.

Scheduling algorithms are used in real-time systems to ensure that every task will meet its deadline. In order to be able to construct schedules in which no tasks miss deadlines, the system must know the maximum computational time needs of each task. In conventional real-time systems, these maximum computational needs are determined at design-time, and the scheduling is also fixed at design-time, so that the system will never miss a deadline.

While this approach is suitable in static systems where normal operating conditions always prevail, in practice several dynamic situations may arise which affect the scheduling. Tasks may overrun their expected computation time due to larger amounts of input data which need processing, or because an iterative algorithm takes a longer time to converge. Faults may occur in the computational environment, reducing the amount of resources available to the computation. This in turn may lead to longer waiting time for resources, and thus cause missed deadlines.

The concept of imprecise and aproximate computations have emerged as the basis of a new approach to dealing with these issues. In dynamic situations, where the time and resources are not enough for computations to complete within the deadline, there may still be enough resources to produce approximate results which may at least avert a catastrophe. This idea is the basis of the imprecise computation approach, which enables real-time programs to produce partial results using less time and resources.

This general concept of using partial results when exact results cannot be produced within the deadline has been used in systems for many years. However, recently this concept has been formalized using the ideas of anytime algorithms and imprecise computation, and specific techniques have been developed for designing programs which can produce partial results, and for developing systems which can support imprecise computation techniques.

The concept of imprecise and approximate computations can be utilized in a wide variety of application areas, including signal processing, machine vision, databases, networking, and dependability in real-time systems. In this volume, we present a collection of papers which describe the current status of the research in this area. The authors of these papers are well-known for their contributions to this field. These papers not only describe research results already obtained, but also discuss directions for further research. Thus, this volume would be of especial interest to researchers who wish to familiarize themselves with the current accomplishments in this area with a view to extending the research results.

This volume is also useful to practitioners interested in building dynamic real-time systems, which must deal safely with resource unavailability. Imprecise and approximate computation has a variety of applications in systems where resource limitations and the requirements of continuous system operation create situations where sometimes computations cannot be carried through to the desired completion. In these cases, the techniques of imprecise computation facilitate the generation of partial results that may enable the system to operate safely and avert catastrophe. The papers in this volume describe some techniques and issues in building systems with these capabilities.

This book may also be used as supplementary reading in a course on real-time systems. The papers contained in this book include a mix of both theoretical papers that present fundamental concepts, and practical papers which discuss issues in building systems and applications. Thus they are well-suited to giving students a balanced perspective of this field.

The first three chapters in this volume discuss techniques for scheduling tasks which allow imprecise computation. The paper by Zhao, Lim, Liu and Alexander discusses the problem of dealing with transient overload situations. The paper by Garvey and Lesser presents a new concept called design-to-time scheduling. The paper by Leung summarizes several other scheduling results in this field.

The next four chapters discuss issues in the design of systems which support imprecise computation. Zilberstein and Russell discuss how utility values may be used in implementing anytime algorithms. The paper by Burns discusses how imprecise computation may be implemented using Ada 9X. Yu and Lin address the use of imprecise computation to provide fault-tolerance, in the context of the Mach operating system. Vrbsky and Liu discuss how imprecise computation may be applied to databases.

The final two chapters in this volume address theoretical foundations. The paper by A. del Val discusses approximation in belief update systems. The paper by Yen and Natarajan presents a decision-theoretic model for analyzing the operation of imprecise computation techniques.

It is our hope that this volume will serve to acquaint the reader with the variety of applications of imprecise and approximate computation, and serve as a guide to both implementation of systems using these techniques, and to avenues for further research in this emerging field.

IMPRECISE AND APPROXIMATE COMPUTATION