FAULT-TOLERANT REAL-TIME SYSTEMS

The Problem of Replica Determinism

THE KLUWER INTERNATIONAL SERIES IN ENGINEERING AND COMPUTER SCIENCE

REAL-TIME SYSTEMS

Consulting Editor
John A. Stankovic

RESPONSIVE COMPUTER SYSTEMS: ***Steps Toward Fault-Tolerant Real-Time Systems,*** by Donald Fussell and Miroslaw Malek, ISBN: 0-7923-9563-8

IMPRECISE AND APPROXIMATE COMPUTATION, by Swaminathan Natarajan, ISBN: 0-7923-9579-4

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

FAULT-TOLERANT REAL-TIME SYSTEMS

The Problem of Replica Determinism

by

Stefan Poledna
Technical University Vienna

Foreword by
H. Kopetz

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

for Hemma

Contents

List of Figures

List of Tables

Foreword

by H. Kopetz

Technical University of Vienna

Hard real-time computer systems, i.e. real-time systems where a failure to meet a deadline can cause catastrophic consequences, are replacing an increasing number of conventional mechanical or hydraulic control systems, particularly in the transportation sector. The vastly expanded functionality of a digital control system makes it possible to implement advanced control algorithms that increase the quality of control far beyond the level that is achievable by a conventional control system.

Computer controlled fly-by-wire technology has been applied widely in the military sector. This technology is now gaining acceptance in the civilian sector as well, such as the digital flight control systems of the Airbus A320 or the Boeing B777 airplane. In these applications, the safety of the plane depends on the reliability of the real-time computer system. A similar development can be observed in the automotive sector. After the successful deployment of computer technology in non safety-critical automotive applications, such as body electronics, the computer control of core vehicle functions, such as engine, brakes, or suspension control is being considered by a number of automotive companies. The benefits that are expected to result from the advanced digital control of core vehicle functions in the automobile are impressive: increased stability and safety of the vehicle, improved fuel efficiency, reduced pollution, etc., that will lead to a safer, more economical, and more comfortable vehicle operation. The mass market of the automotive sector—more than 50 million vehicles are produced worldwide every year—is expected to lead to very cost effective highly integrated computer system solutions that will have a dominant influence on many of the other real-time computer applications. It is therefore expedient to develop new real-time computer system architectures within the constraints given by the automotive applications.

In safety critical applications, such as a drive by wire system, no single point of failure may exist. At present the approach to computer safety in cars is approached at two levels. At the basic level a mechanical system provides the proven safety level that is considered sufficient to operate the car. The computer system provides optimized performance on top of the basic mechanical system. In case the computer system fails cleanly, the mechanical system takes over. Consider, for example, an Anti-lock Braking System (ABS). If the computer fails, the "conventional" mechanical brake system is still operational. In the near future, this approach to safety may reach its limits for two reasons:

(1) If the performance of the computer controlled system is further improved, the "distance" between the performance of the computer controlled system and the performance of the basic mechanical system is further increased. A driver who gets used to the high performance of the computer controlled system might consider the fall-back to the inferior performance of the mechanical system already a safety risk.

(2) The improved price/performance of the microelectronic components will make the implementation of fault-tolerant computer systems cheaper than the implementation of mixed (computer/mechanical) systems. Thus there will be a cost pressure to eliminate the redundant mechanical system.

A solution out of this dilemma is the deployment of a fault-tolerant computer system that will provide the specified service despite a failure of any one of its components. This leads naturally to domain of distributed fault-tolerant hard real-time systems. The stringent timing constraints in many automotive applications—in the millisecond range or below—require the implementation of actively redundant systems. In actively redundant systems the failure in any one of the redundant computational channels is masked immediately by the availability of a set of correct results. A necessary prerequisite for the implementation of active redundancy is the systematic solution of the problem of replica determinism: the assurance that all replicated channel will visit the same computational states at about the same point in time.

The property of replica determinism is also important from the point of view of testability. The sparse time-base of a replica-determinate real-time system makes it possible to specify exactly every test case in the domains of time and value. The reproducibility of the test results, which is a consequence of replica determinism, simplifies the validation of concurrent systems significantly.

The topic of this book, which is a revised version of a Ph.D. Thesis submitted at the Technical University of Vienna, is the systematic treatment of the problem of replica determinism in fault-tolerant real-time systems within the constraints given by the automotive environment. Some of the more formal chapters of the thesis are not included in this work and can be found in the original document. To familiarize the reader with the selected application domain, a special chapter—chapter two—has been introduced that explains the problems and constraints in the field of automotive electronics in some detail.

It was the goal of this research work to find theoretically sound system solutions to the problem of replica determinism that can be implemented within the economic and technical constraints of the automotive industry. This resulted in the formulation of a set of challenging research problems, e.g., the question about the minimal amount of information that has to be exchanged between two replicated computers in order to agree on a single view of the world. The exact specification and systematic analysis of these problems and the presentation of efficient solutions to these problems are a major contribution to the art of designing fault-tolerant hard real-time systems.

We hope that his book will be of great value to designers and implementers of hard real-time computer systems in industry, as well as to students studying the field of distributed fault-tolerant real-time computing.

H. Kopetz

Preface

Real-time computer systems are very often subject to dependability requirements because of their application areas. Fly-by-wire airplane control systems, control of power plants, industrial process control systems and others are required to continue their function despite faults. Therefore, fault-tolerance and real-time requirements constitute a kind of *natural* combination in process control applications. Systematic fault-tolerance is based on redundancy which is used to mask failures of individual components. The problem of replica determinism thereby is to assure that replicated components show consistent behavior in the absence of faults. It might seem trivial that, given an identical sequence of inputs, replicated computer systems will produce consistent outputs. Unfortunately, this is not the case. The problem of replica nondeterminism and the presentation of its possible solutions is the subject of the present work.

The field of automotive electronics is an important application area of fault-tolerant real-time systems. Systems like anti-lock braking, engine control, active suspension or vehicle dynamics control have demanding real-time and fault-tolerance requirements. These requirements have to be met even in the presence of very limited resources since cost is extremely important. Because of its interesting properties this work gives an introduction to the application area of automotive electronics. The requirements of automotive electronics are a topic in the remainder of this work for discussion and are used as a benchmark to evaluate solutions to the problem of replica determinism. The introductory chapter on automotive electronics is self-contained and can be read independently of the remaining chapters.

Following the chapter on automotive electronics a short presentation of the system model and related terminology for fault-tolerant real-time systems is given. This chapter starts the second part of the book which discusses the problem of replica determinism and possible solutions to the problem. First, a generally applicable definition of replica determinism is introduced. Based on this definition possible modes of non-deterministic behavior are then presented. For system design a characterization of *all* the sources of replica non-determinism is important. Such a characterization is given.

The fact that computer systems behave non-deterministically raises the question as to what the appropriate methodologies and implementations for replica determinism enforcement are. This problem is discussed with consideration to different aspects such as communication, synchronization, failures and redundancy preservation. Finally, the problem of replica determinism enforcement is discussed for automotive electronics and systems that have to respond within a short latency period. It is shown that the replication strategies active replication, semi-active replication and passive replication cannot fulfill the given requirements. For that reason two new methodologies are introduced. Firstly, a communication protocol for agreement on external events with a minimum amount of information and, secondly, the concept of timed messages is introduced which allows efficient use of preemptive scheduling in replicated systems. By applying the newly presented methodologies it is shown that replication and systematic fault-tolerance can be used in the area of automotive electronics.

This work is a revised version of my dissertation which was submitted to the Technical University of Vienna in April 1994. Greater emphasis has been especially placed on the application area of automotive electronics. This work has been supported, in part, by the ESPRIT Basic Research Project 'Predictably Dependable Computing Systems' PDCS II.

The valuable support of a number of people have made this work possible. Foremost, I would like to thank my thesis advisor Prof. Herman Kopetz for his many useful suggestions and the interesting discussions which were most fruitful for this work and which served to further my scientific interest. I would also, like to thank him for providing the foreword to this book. Furthermore, my gratitude goes to my colleagues at the Technical University of Vienna and at Robert Bosch. In addition I would like to thank Patrick Henderson and Christopher Temple for their willingness to proofread the manuscripts. Last but not least, I would like to make special mention of the great support and valuable inputs given to me by my friend Ralf Schlatterbeck.

Comments and suggestions concerning the book will be welcomed and can be sent to me by e-mail at stefan@vmars.tuwien.ac.at.

Stefan Poledna