William Fornaciari · Dimitrios Soudris
Editors

Harnessing Performance Variability in Embedded and High-performance Many/Multi-core Platforms

A Cross-layer Approach

Springer

Harnessing Performance Variability in Embedded and High-performance Many/Multi-core Platforms

William Fornaciari • Dimitrios Soudris
Editors

Harnessing Performance Variability in Embedded and High-performance Many/Multi-core Platforms

A Cross-layer Approach

Springer

Editors
William Fornaciari
DEIB – Dipartimento di Elettronica
Informazione e Bioingegneria
Politecnico di Milano
Milano, Italy

Dimitrios Soudris
School of Electrical & Comp. Engineering
Natl. Technical Univ. of Athens
Athens, Greece

ISBN 978-3-319-91961-4 ISBN 978-3-319-91962-1 (eBook)
https://doi.org/10.1007/978-3-319-91962-1

Library of Congress Control Number: 2018960208

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Introduction

This book aims at providing a comprehensive view of the possible solutions to the problem of ensuring dependable performance, by collecting the best practices from the embedded and high-performance worlds.There exists a broad range of possible sources of performance variability, spanning from the hardware (silicon) aspects up to the design of the applications and the management of the resources. To worsen the scenario, it is no longer possible to underestimate the impact of power and thermal management and of process variation, especially in the perspective to deal with reliability and aging issues. The goal of this book is to share the experience maturated during the HARPA project and to complement such information with background concepts and chapters, to create a scientific reference for both academic and leading-edge companies.

The HARPA project's goal has been to enable next-generation embedded and high-performance heterogeneous many-core processors to cost-effectively confront variations and yet provide dependable performance: correct functionality and timing guarantees throughout the expected lifetime of a platform under thermal, power, and energy constraints. The HARPA team is composed of industry and academic partners across Europe specialized in fields covering all the abstraction layers, from hardware to application level. One of the primary objectives of HARPA project is accomplished by means of an accurate yet efficient modeling of performance variability in scaled technologies. Capturing and simulating variability threats is an important step toward providing solutions in mitigating them.

The HARPA solution revolves around a cross-layer approach encompassing all pertinent actors in the system stack. At the core of HARPA, a middleware implements an innovative control engine that steers software/hardware knobs based on information from monitors strategically dispersed in the system. We have explored a large range of such knobs and monitors across different abstraction levels. This engine relies on technology models to identify and exploit various types of platform slack—such as performance, power, energy, temperature, lifetime, and structural (i.e., hardware)—to restore the timing guarantees and ensure expected lifetime in the presence of time-dependent variations. The challenge of dependable performance

is to ensure timing correctness, whereas for high-performance applications it is paramount to ensure load balancing during parallel phases and fast execution. Both are achieved with a combination of a fast control engine with reaction times in the range of a few msec and an OS-level engine reacting slower. These concepts have been demonstrated in three different representative application case studies from the wireless, flood monitoring, and rockfall monitoring domains.

Among different variability phenomena, attention is given to the BTI effects, since they are expected to be the dominant ones. For this purpose, a proper estimation flow that models the dependence of variability at circuit and block level on the operating conditions (temperature and voltage), workload, and technology parameters has been designed and developed. Experimental results have shown running different application benchmarks of interest. The flow is based on commercial ASIC tools and in-house scripts: from the transistor-level representation of the circuit, and the transistor-level (real) workload, BTI variation based on defect-centric models is quantified, and its impact on the critical path delay of the reference circuit is then analyzed through standard STA run.

Although the impact on BTI is driven by the trends in current technology, the system-level perspective is a joint result of different failure mechanisms. To close the gap with system-level analysis, an experimental setup to measure the impacts of stress-induced variability at the system level, from a functional perspective, has been designed and deployed. The proposed setup allows us to directly measure from the hardware the time between successive faults. While modeling and deeply understanding variability impact is an important part of the work, a fundamental objective of the project is to provide solutions to mitigate reliability threats and ensure dependable system performance. Toward this direction, the HARPA engine has been developed, implementing various control frameworks across the system stack. The goal is to exploit different manifestations of platform slack (i.e., slack in performance, power, energy, temperature, lifetime, and structures/components), in order to ascertain timing guarantees throughout the lifetime of the device.

The reader is steered into the concept developed during the project passing through three main stages: background concepts, design methodologies, and analysis of use cases. To simplify the navigation of the book to the reader, the discussion is further split into 6 parts and finally 13 agile chapters.

The beginning of the first part is an introductory chapter providing an overview of the target problems and of the approach to their mitigation developed during the HARPA project, which is a sort of reference guide to move into the book concepts and chapters. The other two chapters of the first part provide a background on the evolution of the computing platforms, with a particular focus on the processor architectures and some foundations on the aging effects from the physics of the phenomenon up to the CAD aspects.

The second part starts entering into the run-time management technologies with two chapters explaining first of all the solutions working at the operating systems level and then how to cope with the problem to achieve a fast and low overhead thermal control of multi-many core processors.

The following two chapters, constituting the third part, move into a finer grain of the design considering the RT level and how to design an ad hoc run-time engine to deal with real-time performance requirements. A comprehensive view of the monitors and knobs is the objective of the fourth part of the book, focusing on both the system-level solutions and on a specific new brick of the emerging platforms, namely, the Network-on-Chip. A coordinated application of the background and concepts presented so far is carried out in the two chapters composing the fifth part of the book. The first of them is affording the problem of properly modeling and simulating BTI-induced degradation in the computing platform and its impact on the timing performance, while the second is the description of a proof of concept performed on a real platform.

The final part of the book figures out how the methodologies and techniques here presented can be effectively applied on real-size problems. To this purpose, two use cases out of the three developed during the HARPA project are presented: Floreon+ to show the possible benefits achievable on application and architectures belonging to the high-performance computing domain and Beesper, which is a typical low-end embedded system submitted to severe reliability and energy management constraints.

This book is probably the first one trying to cover all the critical aspects of the computing continuum and, as we said, is well suited both for engineers working in leading-edge companies and for courses at the master's and PhD level.

Special thanks to Antonio Gonzalez and Jörg Henkel for their availability to cooperate with us in this project, writing two excellent chapters of the book and with the hope to extend in the future the scientific cooperation on new challenges.

We hope you will enjoy reading the book, as the authors enthusiastically worked together for 4 years.

Milano, Italy — William Fornaciari
Athens, Greece — Dimitrios Soudris

Contents