POWER OPTIMIZATION AND SYNTHESIS AT BEHAVIORAL AND SYSTEM LEVELS USING FORMAL METHODS

POWER OPTIMIZATION AND SYNTHESIS AT BEHAVIORAL AND SYSTEM LEVELS USING FORMAL METHODS

by

Jui-Ming Chang
Cadence Design Systems, Inc.

Massoud Pedram
University of Southern California

SPRINGER SCIENCE+BUSINESS MEDIA, LLC

Library of Congress Cataloging-in-Publication Data

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

ISBN 978-1-4613-7368-1 ISBN 978-1-4615-5199-7 (eBook)
DOI 10.1007/978-1-4615-5199-7

Printed on acid-free paper.

This book is dedicated to the development of greater understanding and the reduction of conflicts among people.

Contents

List of Figures

List of Tables

About the Authors

Dr. Jui-Ming Chang is a senior member of the technical staff in Cadence Design Systems, San Jose, CA. He received his B.S. degree in Electrical Engineering from National Taiwan University in 1989 and his M.S. and Ph.D. degrees in computer engineering from the University of Southern California, majoring in VLSI CAD in 1993 and 1998, respectively. His research interests include CAD of VLSI circuits and systems, specializing in behavioral-level (high-level) and system-level (system-on-chip) synthesis. More recently, he has been working on layout synthesis for deep submicron designs, and power estimation and optimization in digital circuits.

Dr. Massoud Pedram is an associate professor of the Department of Electrical Engineering - Systems at the University of Southern California. He received his B.S. degree in Electrical Engineering from California Institute of Technology and M.S. and Ph.D. degrees in Electrical Engineering and Computer Sciences from the University of California, Berkeley in 1989 and 1991, respectively. He is a recipient of the NSF's Young Investigator Award (1994) and the Presidential Faculty Fellows Award (a.k.a. PECASE Award) (1996). His research has received a number of awards including an ICCD Best Paper Award, a DAC Best Paper Award, and an IEEE Transactions on VLSI Systems Best Paper Award.

Dr. Pedram was the co-founder and General Chair of the 1995 International Symposium on Low Power Design and the technical co-chair and general co-chair of the 1996 and 1997 International Symposium on Low Power Electronics and Design, respectively. He has given several tutorials on low power design at major CAD conferences and forums including DAC and ICCAD. He has

published more than 100 journal and conference papers, and co-edited two books on low power design methodologies and techniques. His current research focuses on developing methodologies, techniques, and software tools for lowering the power dissipation in electronic circuits and systems, and on design flows and algorithms for logical/physical codesign of deep submicron circuits. Dr. Pedram is a senior member of IEEE - Circuits and Systems Society and ACM - SIGDA, and an associate editor of the IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems and of the ACM Transactions on Design Automation of Electronic Systems.

Foreword

System-level design is a challenging task because engineers are confronted with a variety of problems, including the need of trading off different objectives, such as performance and power consumption, in the search for design implementations with increasing features, constraints and complexity.

When looking at electronic systems, whether computational or embedded, as information processing machines, an immediate question comes up on the efficiency by which they perform their tasks, as in the case of thermo-dynamical machines. Designing electronic systems with high energy efficiency, or equivalently designing systems that perform under bounds on electric energy consumption, is one of the essential problems to be solved. Despite efforts in this direction in the last decade, new solutions - such as those presented in this book - are crucial for the improvement of system design technologies. Systems with increasingly higher energy efficiency are required by the ongoing needs of batter-power portable systems to support complex software application programs, and by the desire of reducing the environmental impact of all electronic systems.

Design technologies for system-level design have evolved through the years while facing the formidable challenge of addressing computationally complex problems. The widespread use of heuristics in solving design problems has enabled the creation of several computer-aided design tools in a short time. Nevertheless, the quality of the results achieved by design tools based on heuristics is hard to quantify. Algorithms with guaranteed optimality properties are highly-desired because they can enable a thorough search of the design space.

This book makes an important contribution to the field of system design technologies by presenting a set of algorithms with guaranteed optimality properties, that can be readily applied to system-level design. This contribution is timely, because it fills the need of new methods for a new design tool generation, which supports the design of electronic systems with even more demanding requirements.

Giovanni De Micheli, Professor
Stanford University

Preface

With the Moore's law still in effect, integrated circuit densities and operating speeds continue to rise at an exponentail rate. Chips however cannot get larger and faster without a sharp increase in power consumption beyond the current levels. Minimization of power consumption in VLSI chips has thus become an important design objective. In fact, with the explosive growth in demand for portable electronics and the usual push toward more complex functionality and higher performance, power consumption has in many cases become the limiting factor in satisfying the market demand.

The low power design challenge has been met by an active research and development community both in industry and academia. Rapid advances are taking place in low-power process technologies, architecture and circuit optimization techniques, power-aware simple and complex cell design, use of variable and/or multiple supply voltages and dynamic power management schemes, and low power computer aided design (CAD) tools from system and software levels to layout and transistor levels. In particular, a new generation of power-conscious CAD tools are coming into the market to help designers estimate, optimize and verify power consumption levels at most stages of the IC design process. These tools are especially prevalent at the register-transfer level and below. There is great need for similar tools and capabilities at the behavioral and system levels of the design process.

The increased degree of automation of industrial design frameworks has produced a substantial change in the way digital ICs are developed. The design of modern systems usually starts from specifications given at a very high level of abstraction. This is because existing EDA tools are able to automatically pro-

duce low-level design implementations directly from descriptions of this type. Circuit and system designers need tools that allow them to explicitly control the power budget during the early phases of the design process. This is because the power savings obtainable through automatic optimization early in the design process is usually more significant than that achievable by means of lower level optimization.

This need has not gone unnoticed. Many researchers and CAD tool developers are working on high-level power modeling and estimation, as well as power-constrained high-level synthesis and optimization. Techniques and tools alone are however insufficient to optimize the VLSI circuit power dissipation – a consistent and convergent design methodology is required as well. Components of such a low power design methodology include: upfront specification, early analysis and optimization, forward timing and power constraint propagation and backward capacitance annotation, as well as multi-level power simulation and verification. Tools that support such a flow would include power macro-models for such library primitives as gates, adders, multipliers and register files as well as more complex functions such as memories, controllers, encoder/decoders and in general, intellectual property (IP) blocks.

The present book is written to address some of the key problems in power analysis and optimization early in the design process. In particular, this book focuses on power macro-modeling based on regression analysis and power minimization through behavioral transformations, scheduling, resource assignment and hardware/software partitioning and mapping. What differentiates this book from the other published work on the subject is that the mathematical basis and formalism behind our algorithms and the optimality of the these algorithms subject to the stated assumptions.

This book is organized into six chapters. Chapter 1 lays the groundwork by presenting an overview of the behavioral-level and system-level synthesis techniques and design representations. In Chapters 2 and 3, we present optimal algorithms for activity-driven register and module allocation and binding. The power consumed by a resource mainly depends on the input switching activities induced by the data being stored or processed. Since the patterns flowing through a circuit may have specific probability distribution, the way registers and modules are allocated in a control/data flow graph may heavily impact the switching activities at the interface of the resources. These chapters present

graph-theory-based algorithms for power-efficient allocation and binding of registers and modules based on accurate computation of the probability density functions at the inputs of various resources, given probability distributions for the system primary input. In Chapter 4, we describe an algorithm based on dynamic programming for solving the multiple-voltage scheduling problem. The technique, which is based on dynamic programming requires the availability of accurate timing and power models for the modules in the RTL library for all possible supply voltage levels. Using this information, and by a post-order traversal of the data flow graph, it is then possible to calculate the energy-delay trade-off curves at each node of the graph. Supply voltage level assignment and scheduling of all operations take place during a pre-order traversal of the data flow graph. In Chapter 5, we provide an optimal algorithm for finding the minimum cost (area and/or energy dissipation) hardware/software partitioning, and mapping of a generalized task graph subject to performance constraints. Processes in a generalized task graph may communicate with each other at times other than the beginning or end of their lifetimes by various blocking/non-blocking communication mechanisms. The coarse-grain HW/SW partitioning/mapping technique used a modified (symbolic) dynamic programming algorithm with binning strings to produce a globally optimal solution subject to model assumptions. When needed, we also present power estimation models and techniques to support the optimization algorithms. Chapter 6 presents the book summary and outline of possible research venues.

We wish to express our appreciation to the students in the low power CAD group at USC for providing a stimulating and challenging environment, to Professor Giovanni De Micheli for writing the foreword for the book, to Carl Harris and staff at Kluwer Academic Publishers for their help in editing and publishing the book, and finally to our families for their indulgence of our interest in this project.

Jui-Ming (Raymond) Chang, San Jose
Massoud Pedram, Los Angeles