

The background of the cover is a detailed, high-contrast image of a circuit board. It features a complex network of gold-colored traces and pads on a dark, possibly black or dark green, substrate. The traces are arranged in a dense, interconnected pattern, typical of modern electronic packaging. The lighting is dramatic, with some areas appearing brighter than others, creating a sense of depth and highlighting the metallic texture of the traces.

Unique Chips and Systems

**Edited by
Eugene John
Juan Rubio**



CRC Press
Taylor & Francis Group

TN402
U58

Unique Chips and Systems

Edited by
**Eugene John
Juan Rubio**



E2009003547



CRC Press

Taylor & Francis Group

Boca Raton London New York

CRC Press is an imprint of the
Taylor & Francis Group, an **informa** business

CRC Press
Taylor & Francis Group
6000 Broken Sound Parkway NW, Suite 300
Boca Raton, FL 33487-2742

© 2008 by Taylor & Francis Group, LLC
CRC Press is an imprint of Taylor & Francis Group, an Informa business

No claim to original U.S. Government works
Printed in the United States of America on acid-free paper
10 9 8 7 6 5 4 3 2 1

International Standard Book Number-13: 978-1-4200-5174-2 (Hardcover)

This book contains information obtained from authentic and highly regarded sources. Reprinted material is quoted with permission, and sources are indicated. A wide variety of references are listed. Reasonable efforts have been made to publish reliable data and information, but the author and the publisher cannot assume responsibility for the validity of all materials or for the consequences of their use.

Except as permitted under U.S. Copyright Law, no part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, microfilming, and recording, or in any information storage or retrieval system, without written permission from the publishers.

For permission to photocopy or use material electronically from this work, please access www.copyright.com (<http://www.copyright.com/>) or contact the Copyright Clearance Center, Inc. (CCC) 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. CCC is a not-for-profit organization that provides licenses and registration for a variety of users. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

Trademark Notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Library of Congress Cataloging-in-Publication Data

John, Eugene.
Unique chips and systems / Eugene John and Juan Rubio.
p. cm.
Includes bibliographical references and index.
ISBN 978-1-4200-5174-2 (hbk. : alk. paper)
1. Systems on a chip. 2. Computer networks--Equipment and supplies. I. Rubio, Juan, 1973- II. Title.

TK7895.E42J64 2007
621.3815--dc22

2007024835

Visit the Taylor & Francis Web site at
<http://www.taylorandfrancis.com>

and the CRC Press Web site at
<http://www.crcpress.com>

Unique Chips and Systems

Computer Engineering Series

Series Editor: Vojin Oklobdzija

*Coding and Signal Processing for
Magnetic Recording Systems*

Edited by Bane Vasic and Erozan M. Kurtas

*Digital Image Sequence Processing,
Compression, and Analysis*

Edited by Todd R. Reed

Low-Power Electronics Design

Edited by Christian Piguet

Unique Chips and Systems

Edited by Eugene John and Juan Rubio

Preface

Integrated circuits are the enabling technology for the modern information age. Advanced systems are built using state-of-the-art semiconductor chips. Computing, communication, and network chips fuel the information technology era. The demands of emerging software applications can be met only with unique chips and systems. The integration ability presented by modern semiconductor technology presents opportunities; however, the requirements posed by power consumption, reliability, and form factor present challenges. This book presents fourteen chapters dealing with several systems and chips that present unique approaches to designing future computing and communication chips and systems.

Chapter 1 presents the TRIPS processor architecture and microarchitecture. TRIPS is a unique architecture that seeks to better exploit uniprocessor-level concurrency by changing the way instruction-level concurrency is expressed to the hardware, thereby extending the scaling of uniprocessors and enabling more efficient multiprocessors. TRIPS uses an explicit data graph execution (EDGE) instruction set architecture to efficiently encode concurrency in its dataflow execution model. The TRIPS microarchitecture uses a distributed, tiled microarchitecture that supports dynamic out-of-order execution. It is partitioned for scalability and implements deep speculation and latency tolerance.

Chapter 2 describes the Centaur Technology x86 processor with several data security features. Centaur Technology (a part of VIA Technologies Inc.) integrated several security features into the x86 processor, with little increase in die size or development effort. The chapter presents the hardware security features, and describes the implementation of the AES encryption hardware, the secure hash algorithm (SHA) hardware and the Montgomery multiplier—all aimed at improving the security of the processor.

Chapter 3 presents the ARM Cortex-A8 processor, a sub-1 watt processor that provides high performance for general purpose and media applications. The processor performs superscalar execution; yet, it is designed to be energy efficient. The microarchitecture, machine efficiency, and operating frequency are decided with energy efficiency as a primary criterion. Multimedia and graphics applications are supported with a 64-bit SIMD unit.

Chapter 4 presents a highly parallel signal processor, the RACE-Hypercube processor, which achieves up to 1 trillion bytes/sec at a relatively low clock frequency of 250 MHz. The processor allows the selection of a variety of configuration parameters.

Chapter 5 presents an asynchronous FPGA design—the RASTER architecture. The challenges and limitation of a clocked design are overcome with a self-timed (asynchronous) design, resulting in higher performance per watt. The RASTER architecture consists of an FPGA logic cell that uses a unique method of intercell communication. Simulation shows data throughput rates of up to 1.3 GHz at the 90nm process on a benchmarking suite of small FPGA designs.

Chapter 6 presents another unique chip—the continuation-based Fuce multithreading processor. The Fuce processor from Kyushu University, Japan, is based on the dataflow computing model. The Fuce processor pursues parallel execution of threads with high parallel processing and compatibility. Fuce means “fusion of communication and execution.” The Fuce processor executes multiple threads using the exclusive multithread execution model, which is derived from dataflow computing. The Fuce processor aims to fuse the interprocessor execution and interprocessor communication. The Fuce processor unifies processing inside the processor and communication with external processors using events and threads.

Chapter 7 is a study of a processor with dual thread execution modes. The authors present the use of additional cores on a processor for two purposes: (1) to execute subordinate threads, and (2) to execute speculative threads. Threads are spawned to the available processing cores to exploit thread-level parallelism. Performance analysis using SPEC CPU2000 benchmarks show higher improvement using subordinate threads rather than speculative threads. A processor that can switch execution modes between the two approaches is also investigated since many applications alternate between different types of phases during their execution. Such an adaptive processor is seen to be 17 percent better than the subordinate thread mechanism alone.

Adaptive power management of computer systems has become extremely important in recent years. Such techniques heavily rely on variation of power during execution of applications. Chapter 8 presents power phases in commercial and scientific workloads running on enterprise-class hardware. Power consumption of CPU, I/O, and disk subsystems is measured using power sensors and phase behavior of applications is studied.

Future chips are driven by emerging and future applications. A workload that is most demanding of computational power and speed is computer graphics and visualization. Gaming has driven this quest for function and speed to such a point that graphics chips, independent of the driving computer system, have more gates than the latest CPU and many times the arithmetic power. And yet, there are aspects of graphics that still overly consume the power of systems. In Chapter 9, example graphics applications that need enormous computing power are presented. The author seeks to provide compact geometric representations of shapes so that rendering (displaying on the screen) can be more efficiently performed. He shows a close relationship between quadratic Bezier curves (QBCs) and iterated function systems (IFSs) to manipulate 2D sets that resemble 3D sets in the real world. He also

demonstrates the value of segmenting 3D triangle meshes that represent human teeth, thus dramatically accelerating visualization processes.

In Chapter 10, the authors illustrate the use of hardware accelerators built from field programmable gate arrays (FPGAs), graphic processing units (GPUs), or SIMD processor arrays for high performance computing. Such a system can be considered as a two-level processing system, consisting of the conventional processing nodes and the acceleration hardware connected over a high-speed network. In this chapter, researchers from the Los Alamos National Laboratory describe the use of such systems for a class of applications that use wavefront algorithms. These algorithms are characterized by a specific order in which cells are processed. The improvement in performance from accelerators such as the Clearspeed CSX600 SIMD accelerator is presented.

In Chapter 11, characteristics of a bioinformatics application are presented. Computational biology has become an important workload for high performance computers. Multiple-sequence alignment applications are important bioinformatics applications. Twelve multiple sequence alignment programs with a variety of alignment approaches are analyzed for performance of the cache, trace cache, branch predictor, phase behavior, and so on.

Embedded systems are inherently real time systems—they must control and compute as demanded by events. And the larger systems they are part of may demand a significant number of parallel processes going; for example, the most lavishly outfitted BMW automobile has an excess of 100 microcontrollers in charge of its many operations. Ravenscar is a subset of the Ada programming language designed for real-time computing. In Chapter 12, the authors present a Ravenscar, hardware-implemented run-time kernel with delay queues that allows for accurate analysis of application timing behavior. Formal state models and their simulations as well as hardware implementation are presented. The authors describe the corresponding VHDL state machines and demonstrate that the required levels of parallelism, hardware requirements, and timing granularity can be achieved.

In Chapter 13, an error correction scheme for a network-on-chip (NOC) is presented. The increased susceptibility of on-chip networks to various sources of error necessitates strategies to handle errors. A forward error correction scheme employing a low density parity check code (LDPC) is presented in this chapter. The presented LDPC is a linear block code suitable for low latency, high gain, and low power design because of its streamlined forward-only data flow structure.

Chapter 14 presents silicon-based on-chip optical interconnects and their use in reducing thermal constraints in a high performance clustered multi-threaded processor. Increased integration in modern semiconductor technologies often results in regions of the chip with very high power densities or hot spots. One technique to reduce the thermal concerns from the hot spots is to intermix hot and cold units, however, at the cost of increasing communication distances between blocks. Silicon-based optical interconnects are shown to

be very valuable for global communication paths in such chips. A significant reduction in thermal constraints without reducing performance is shown in connecting the common front-end with the distributed back-end of a clustered multithreaded processor.

We hope that the readers of this book enjoy the variety of unique systems and chips presented. Most of the chapters in this book are revised versions of selected papers presented at the first, second, and third Workshop on Unique Chips and Systems (UCAS). The first and second UCAS workshops were held in March 2005 and March 2006 in Austin, Texas, and the third UCAS workshop was held in San Jose, California, in April 2007. We would like to thank the authors of the chapters for their contributions. We also wish to thank all those who helped in the process, especially Nora Konopka and Jessica Vakili at CRC Press/Francis & Taylor.

Eugene John

University of Texas at San Antonio

Juan Rubio

IBM Austin Research Laboratory

Editors

Eugene John is a professor in the department of electrical and computer engineering at the University of Texas at San Antonio. He received his Ph.D. in electrical engineering from Pennsylvania State University in 1995. His current research interests include low power circuits and systems, VLSI design, power estimation and optimization, multimedia and network processors, computer architecture, performance evaluation, and biometrics. He is a senior member of the IEEE, IEEE Computer Society, and IEEE Circuits and Systems Society. He is also a member of Eta Kappa Nu, Tau Beta Pi, and Phi Kappa Phi Honor societies.

Juan Rubio is a research staff member at the IBM Austin Research Lab. His interests include computer system architecture, performance analysis, and control systems. At IBM, he explores techniques to model, monitor and manage power; temperature and performance in computer servers; and data centers. His research contributed directly to the development of PowerExecutive™ and EnergyScale™ technology used in IBM systems.

Rubio received a B.S. in electrical engineering from Universidad Santa Maria La Antigua, Panama, in 1997, and an M.S. and Ph.D. in computer engineering from the University of Texas at Austin in 2004.

Contributors

David H. Albonesi Cornell University, Ithaca, New York

Makoto Amamiya Kyushu University, Fukuoka, Japan

Satoshi Amamiya Kyushu University, Fukuoka, Japan

Frank Barry Onward Communications Inc. and Appalachian State University, Boone, North Carolina

Praveen Bhojwani Texas A&M University, College Station, Texas

W. Lloyd Bircher University of Texas at Austin, Texas

Gregory Briggs University of Rochester, Rochester, New York

Doug Burger University of Texas at Austin, Texas

Guoqing Chen University of Rochester, Rochester, New York

Hui Chen University of Rochester, Rochester, New York

Gwan Choi Texas A&M University, College Station, Texas

Tom Crispin Centaur Technology Inc., Austin, Texas

Rajagopalan Desikan University of Texas at Austin, Texas

Saurabh Drolia University of Texas at Austin, Texas

Philippe M. Fauchet University of Rochester, Rochester, New York

Manoj Franklin University of Maryland, College Park, Maryland

Eby G. Friedman University of Rochester, Rochester, New York

Brian C. Gaide University of Texas at Austin, Texas

M. S. Govindan University of Texas at Austin, Texas

Paul Gratz University of Texas at Austin, Texas

Divya Gulati University of Texas at Austin, Texas

Heather Hanson University of Texas at Austin, Texas

Mikhail Haurylau University of Rochester, Rochester, New York

G. Glenn Henry Centaur Technology Inc., Austin, Texas

Adolfy Hoisie Los Alamos National Laboratory, Los Alamos, New Mexico

Masaaki Izumi Kyushu University, Fukuoka, Japan

Chand T. John Stanford University, Stanford, California

Lizy Kurian John University of Texas at Austin, Texas

Stephen W. Keckler University of Texas at Austin, Texas

Darren J. Kerbyson Los Alamos National Laboratory, Los Alamos, New Mexico

Changkyu Kim University of Texas at Austin, Texas

Tao Li University of Florida, Gainesville, Florida

Haiming Liu University of Texas at Austin, Texas

Kristina Lundqvist Massachusetts Institute of Technology, Cambridge, Massachusetts

Rabi Mahapatra Texas A&M University, College Station, Texas

Rania Mameesh University of Maryland, College Park, Maryland

Takanori Matsuzaki Kyushu University, Fukuoka, Japan

Robert McDonald University of Texas at Austin, Texas

Ramadass Nagarajan University of Texas at Austin, Texas

Nicholas Nelson University of Rochester, Rochester, New York

Terry Parks Centaur Technology Inc., Austin, Texas

Gerald G. Pechanek Lightning Hawk Consulting Inc. Cary, North Carolina, and Priest & Goldstein, PLLC, Durham, North Carolina

Nikos Pitsianis Duke University, Durham, North Carolina

Nitya Ranganathan University of Texas at Austin, Texas

Karthikeyan Sankaralingam University of Texas at Austin, Texas

Simha Sethumadhavan University of Texas at Austin, Texas

Sadia Sharif University of Texas at Austin, Texas

Premkishore Shivakumar University of Texas at Austin, Texas

Rohit Singhal Texas A&M University, College Station, Texas

Mihailo Stojancic ViCore Technologies Inc., Palo Alto, California

David Williamson ARM Inc., Austin, Texas

Contents

Preface	vii
Editors	xi
Contributors	xiii
1 Architecture and Implementation of the TRIPS Processor	1
<i>Stephen W. Keckler, Doug Burger, Karthikeyan Sankaralingam, Ramadass Nagarajan, Robert McDonald, Rajagopalan Desikan, Saurabh Drolia, M. S. Govindan, Paul Gratz, Divya Gulati, Heather Hanson, Changkyu Kim, Haiming Liu, Nitya Ranganathan, Simha Sethumadhavan, Sadia Sharif, and Premkishore Shivakumar</i>	
2 High-Performance Data Security in an x86 Processor	41
<i>G. Glenn Henry, Terry Parks, and Tom Crispin</i>	
3 ARM Cortex-A8: A High-Performance Processor for Low-Power Applications	79
<i>David Williamson</i>	
4 A Rotated Array Clustered Extended Hypercube Processor: The RACE-H™ Processor	107
<i>Gerald G. Pechanek, Mihailo Stojancic, Frank Barry, and Nikos Pitsianis</i>	
5 A High-Throughput Self-Timed FPGA Core Architecture	125
<i>Brian C. Gaide and Lizy Kurian John</i>	
6 The Continuation-Based Multithreading Processor: Fuce	177
<i>Masaaki Izumi, Satoshi Amamiya, Takanori Matsuzaki, and Makoto Amamiya</i>	
7 A Study of a Processor with Dual Thread Execution Modes	197
<i>Rania Mameesh and Manoj Franklin</i>	

8 Measurement-Based Power Phase Analysis 217
W. Lloyd Bircher and Lizy Kurian John

**9 Visualization by Subdivision: Two Applications
for Future Graphics Platforms 239**
Chand T. John

**10 A Performance Analysis of Two-Level Heterogeneous
Processing Systems on Wavefront Algorithms 259**
Darren J. Kerbyson and Adolffy Hoisie

**11 Microarchitectural Characteristics and Implications
of Alignment of Multiple Bioinformatics Sequences..... 281**
Tao Li

**12 Towards System-Level Fault-Tolerance Using Formal
Methods and SoC Methodologies 299**
Kristina Lundqvist

**13 Forward Error Correction for On-Chip Interconnection
Networks..... 325**
Praveen Bhojwani, Rohit Singhal, Gwan Choi, and Rabi Mahapatra

**14 Alleviating Thermal Constraints while Maintaining
Performance via Silicon-Based On-Chip
Optical Interconnects..... 339**
*Nicholas Nelson, Gregory Briggs, Mikhail Haurylau, Guoqing Chen,
Hui Chen, Eby G. Friedman, Philippe M. Fauchet, and David H. Albonesi*

Index..... 357

1

Architecture and Implementation of the TRIPS Processor

**Stephen W. Keckler, Doug Burger, Karthikeyan Sankaralingam,
Ramadass Nagarajan, Robert McDonald, Rajagopalan Desikan,
Saurabh Drolia, M. S. Govindan, Paul Gratz, Divya Gulati, Heather
Hanson, Changkyu Kim, Haiming Liu, Nitya Ranganathan, Simha
Sethumadhavan, Sadia Sharif, and Premkishore Shivakumar**

The University of Texas at Austin

CONTENTS

1.1	Introduction.....	2
1.2	ISA Support for Distributed Execution	4
1.2.1	TRIPS Blocks	4
1.2.2	TRIPS Instruction Formats.....	5
1.2.3	Code Generation	6
1.3	A Distributed Microarchitecture.....	8
1.3.1	Global Control Tile (GT)	9
1.3.1.1	Fetch Unit.....	10
1.3.1.2	Refill Unit.....	12
1.3.1.3	Retire Unit	12
1.3.1.4	Next Block Predictor	13
1.3.2	Instruction Tile (IT)	14
1.3.3	Register Tile (RT)	14
1.3.4	Execution Tile (ET)	15
1.3.5	Data Tile (DT)	16
1.3.5.1	Load Processing.....	16
1.3.5.2	Store Processing.....	18
1.3.5.3	Store Tracking	19
1.3.5.4	Memory-Side Dependence Processing.....	19
1.3.5.5	Load/Store Queues	19
1.3.6	Secondary Memory System.....	20
1.3.6.1	OCN Router.....	20
1.3.6.2	Network Address Translation	21
1.4	Distributed Microarchitectural Protocols	22

- 1.4.1 Block Fetch Protocol 22
- 1.4.2 Distributed Execution 24
- 1.4.3 Block/Pipeline Flush Protocol 26
- 1.4.4 Block Commit Protocol 26
- 1.5 Physical Design/Performance Overheads 27
 - 1.5.1 Chip Specifications 27
 - 1.5.2 Chip Verification 29
 - 1.5.3 TRIPS System 30
 - 1.5.4 Area Overheads of Distributed Design 30
 - 1.5.5 Timing Overheads 31
 - 1.5.6 Performance Overheads 32
 - 1.5.6.1 Distributed Protocol Overheads 32
 - 1.5.6.2 Total Performance 34
- 1.6 Related Work 35
 - 1.6.1 Tiled Architectures 35
 - 1.6.2 Dataflow Architectures 36
 - 1.6.3 Superscalar Architectures 36
 - 1.6.4 VLIW Architectures 36
- 1.7 Conclusions 36
- Acknowledgments 37
- References 38



1.1 Introduction

Growing on-chip wire delays, coupled with complexity and power limitations, have placed severe constraints on the issue-width scaling of centralized superscalar architectures. As a result, recent microprocessor designs have backed away from powerful uniprocessors, instead favoring multiple simpler cores on a single die. Partitioning the chip into a collection of processors communicating via a common memory system mitigates some of the technology scaling challenges, but increases the burden on software to provide multiple threads to execute concurrently across the cores.

An alternative is to pursue more powerful uniprocessors, but design them so that they are scalable and tolerant of technology and complexity scaling. Ideally, such wide-issue processors would be *tiled* [30], meaning composed of multiple replicated, communicating design blocks. Because of multicycle communication delays across these large processors, control must be distributed across the tiles. We advocate the use of microarchitectural networks (or *micronets*) for routing control and data among the tiles. Micronets provide high-bandwidth, flow-controlled transport for control or data in a wire-dominated processor by connecting the multiple tiles, each of which is a client on one or more micronets. Higher-level microarchitectural protocols direct global control across the micronets and tiles in a manner invisible to software.