Closing The Gap Between Asic And Custom Tools And Techniques For High Performance Asic Design Author David Chinnery Jun 2002

Reuse Methodology Manual for System-on-a-Chip Designs, Third Edition outlines a set of best practices for creating reusable designs for use in an SoC design methodology. These practices are based on the authors' experience in developing reusable designs, as well as the experience of design teams in many companies around the world. Silicon and tool technologies move so quickly that many of the details of design-for-reuse will undoubtedly continue to evolve over time. But the fundamental aspects of the methodology described in this book have become widely adopted and are likely to form the foundation of chip design for some time to come. Development methodology necessarily differs between system designers and processor designers, as well as between DSP developers and chipset developers. However, there is a common set of problems facing everyone who is designing complex chips. In response to these problems, design teams have adopted a block-based design approach that emphasizes design reuse. Reusing macros (sometimes called "cores") that have already been designed and verified helps to address all of the problems above. However, in adopting reuse-based design, design teams have run into a significant problem. Reusing blocks that have not been explicitly designed for reuse has often provided little or no benefit to the team. The effort to integrate a pre-existing block into new designs can become prohibitively high, if the block does not provide the right views, the right documentation, and the right functionality. From this experience, design teams have realized that reuse-based design requires an explicit methodology for developing reusable macros that are easy to integrate into SoC designs. This manual focuses on describing these techniques. Features of the Third Edition: Up to date; State of the art; Reuse as a solution for circuit designers; A chronicle of "best practices"; All chapters updated and revised; Generic guidelines - non tool specific; Emphasis on hard IP and physical design.

In response to tremendous growth and new technologies in the semiconductor industry, this volume is organized into five, information-rich sections. Digital Design and Fabrication surveys the latest advances in computer architecture and design as well as the technologies used to manufacture and test them. Featuring contributions from leading experts, the book also includes a new section on memory and storage in addition to a new chapter on nonvolatile memory technologies. Developing advanced concepts, this sharply focused book— Describes new technologies that have become driving factors for the electronic industry Includes new information on semiconductor memory circuits, whose development best illustrates the phenomenal progress encountered by the fabrication and technology sector Contains a section dedicated to issues related to system power consumption Describes reliability and testability of computer systems Pinpoints trends and state-of-the-art advances in fabrication and CMOS technologies Describes performance evaluation measures, which are the bottom line from the user’s point of view Discusses design techniques used to create modern computer systems, including high-speed computer arithmetic and high-frequency design, timing and clocking, and PLL and DLL design

Design for Manufacturability and Statistical Design: A Comprehensive Approach presents a comprehensive overview of methods that need to be mastered in understanding state-of-the-art design for manufacturability and statistical design methodologies. Broadly, design for manufacturability is a set of techniques that attempt to fix the systematic sources of variability, such as those due to photolithography and CMP. Statistical design, on the other hand, deals with the random sources of variability. Both paradigms operate within a common framework, and their joint comprehensive treatment is one of the objectives of this book and an important differentation.

This book contains all the topics of importance to the low power designer. It first lays the foundation and then goes on to detail the design process. The book also discusses such special topics as power management and modal design, ultra low power, and low power design methodology and flows. In addition, coverage includes projections of the future and case studies.

Covers the methodology and state-of-the-art techniques of constrained verification, which is new and popular. It relates constrained verification with the also-hot technology called assertion-based design. Discussed and clarifies language issues, critical to both the above, which will help the implementation of these languages.

This book constitutes the refereed proceedings of the Cryptographers' Track at the RSA Conference 2012, CT-RSA 2012, held in San Francisco, CA, USA, in February/March 2012. The 26 revised full papers presented were carefully reviewed and selected from 113 submissions. The papers are organized in topical sections on side channel attacks, digital signatures, public-key encryption, cryptographic protocols, secure implementation methods, symmetric key primitives, and secure multiparty computation.

This book tackles head-on the challenges of digital design in the era of billion-transistor SoCs. It discusses fundamental design concepts in design and coding required to produce robust, functionally correct designs. It also provides specific techniques for measuring and minimizing complexity in RTL code. Finally, it discusses the tradeoff between RTL and high-level (C-based) design and how tools and languages must progress to address the needs of tomorrow’s SoC designs.

While battery capacity is often insufficient to keep up with the power-demanding features of the latest mobile devices, powering the functional advancement of wireless devices requires a revolution in the concept of battery life and recharge capability. Future handheld devices and wireless networks should be able to recharge themselves automaticall

"This reference is a broad, multi-volume collection of the best recent works published under the umbrella of computer engineering, including perspectives on the fundamental aspects, tools and technologies, methods and design, applications, managerial impact, social/behavioral perspectives, critical issues, and emerging trends in the field"--Provided by publisher.

Discover innovative tools that pave the way from circuit and physical design to fabrication processing Nano-CMOS Design for Manufacturability examines the challenges that design engineers face in the nano-scaled era, such as exacerbated effects and the proven design for manufacturability (DFM) methodology in the midst of increasing variability and design process interactions. In addition to discussing the difficulties brought on by the continued dimensional scaling in conformance with Moore's law, the authors also tackle complex issues in the design process to overcome the difficulties, including the use of a functional first silicon to support a predictable product ramp. Moreover, they introduce several emerging concepts, including stress proximity effects, contour-based extraction, and design process interactions. This book is the sequel to Nano-CMOS Circuit and Physical Design, taking design to technology nodes beyond 65nm geometries. It is divided into three parts: Part One, Newly Exacerbated Effects, introduces the newly exacerbated effects that require designers' attention, beginning with a discussion of the lithography aspects of DFM, followed by the impact of layout on transistor performance Part Two, Design Solutions, examines how to mitigate the impact of process effects, discussing the methodology needed to make sub-wavelength patterning technology work in manufacturing, as well as design solutions to deal with signal, power integrity, WELL, stress proximity effects, and process variability Part Three, The Road to DFM, describes new tools needed to support DFM efforts, including an auto-correction tool capable of fixing the layout of cells with multiple optimization goals, followed by a look ahead into the future of DFM Throughout the book, real-world examples simplify complex concepts, helping readers see how they can successfully

handle projects on Nano-CMOS nodes. It provides a bridge that allows engineers to go from physical and circuit design to fabrication processing and, in short, make designs that are not only functional, but that also meet power and performance goals within the design schedule.

This book constitutes the refereed proceedings of the 8th International Workshop on Cryptographic Hardware and Embedded Systems, CHES 2006, held in Yokohama, Japan in October 2006. The 32 revised full papers presented together with three invited talks were carefully reviewed and selected from 112 submissions.

This book is the second of two volumes addressing the design challenges associated with new generations of semiconductor technology. The various chapters are compiled from tutorials presented at workshops in recent years by prominent authors from all over the world. Technology, productivity and quality are the main aspects under consideration to establish the major requirements for the design and test of upcoming systems on a chip.

by Kurt Keutzer Those looking for a quick overview of the book should fast-forward to the Introduction in Chapter 1. What follows is a personal account of the creation of this book. The challenge from Earl Killian, formerly an architect of the MIPS processors and at that time Chief Architect at Tensilica, was to explain the significant performance gap between ASICs and custom circuits designed in the same process generation. The relevance of the challenge was amplified shortly thereafter by Andy Bechtolsheim, founder of Sun Microsystems and ubiquitous investor in the EDA industry. At a dinner talk at the 1999 International Symposium on Physical Design, Andy stated that the greatest near-term opportunity in CAD was to develop tools to bring the performance of ASIC circuits closer to that of custom designs. There seemed to be some synchronicity that two individuals so different in concern and character would be pre-occupied with the same problem. Intrigued by Earl and Andy’s comments, the game was afoot. Earl Killian and other veterans of microprocessor design were helpful with clues as to the sources of the performance discrepancy: layout, circuit design, clocking methodology, and dynamic logic. I soon realized that I needed help in tracking down clues. Only at a wonderful institution like the University of California at Berkeley could I so easily commandeer an ab- bodied graduate student like David Chinnery with a knowledge of architecture, circuits, computer-aided design and algorithms.

This book describes a new design approach for energy-efficient, Domain-Specific Instruction set Processor (DSIP) architectures for the wireless baseband domain. The innovative techniques presented enable co-design of algorithms, architectures and technology, for efficient implementation of the most advanced technologies. To demonstrate the feasibility of the author’s design approach, case studies are included for crucial functionality of advanced wireless systems with increased computational performance, flexibility and reusability. Designers using this approach will benefit from reduced development/product costs and greater scalability to future process technology nodes.

This volume features the refereed proceedings of the 17th International Workshop on Power and Timing Modeling, Optimization and Simulation. Papers cover high level design, low power design techniques, low power analog circuits, statistical static timing analysis, power modeling and optimization, low power routing optimization, security and asynchronous design, low power applications, modeling and optimization, and more.

Explains how to use low power design in an automated design flow, and examine the design time and performance trade-offs Includes the latest tools and techniques for low power design applied in an ASIC design flow Focuses on low power in an automated design methodology, a much neglected area

Power consumption becomes the most important design goal in a wide range of electronic systems. There are two driving forces towards this trend: continuing device scaling and ever increasing demand of higher computing power. First, device scaling continues to satisfy Moore’s law via a conventional way of scaling (More Moore) and a new way of exploiting the vertical integration (More than Moore). Second, mobile and IT convergence requires more computing power on the silicon chip than ever. Cell phones are now evolving towards mobile PC. PCs and data centers are becoming commodities in house and a must in industry. Both supply enabled by device scaling and demand triggered by the convergence trend realize more computation on chip (via multi-core, integration of diverse functionalities on mobile SoCs, etc.) and finally more power consumption incurring power-related issues and constraints. Energy-Aware System Design: Algorithms and Architectures provides state-of-the-art ideas for low power design methods from circuit, architecture to software level and offers design case studies in three fast growing areas of mobile storage, biomedical and security. Important topics and features: - Describes very recent advanced issues and methods for energy-aware design at each design level from circuit and architecture to algorithm level, and also covering important blocks including low power main memory subsystem and on-chip network at architecture level - Explains efficient power conversion and delivery which is becoming important as heterogeneous power sources are adopted for digital and non-digital parts - Investigates 3D die stacking emphasizing temperature awareness for better perspective on energy efficiency - Presents three practical energy-aware design case studies; novel storage device (e.g., solid state disk), biomedical electronics (e.g., cochlear and retina implants), and wireless surveillance camera systems. Researchers and engineers in the field of hardware and software design will find this book an excellent starting point to catch up with the state-of-the-art ideas of low power design.

Here is an extremely useful book that provides insight into a number of different flavors of processor architectures and their design, software tool generation, implementation, and verification. After a brief introduction to processor architectures and how processor designers have sometimes failed to deliver what was expected, the authors introduce a generic flow for embedded on-chip processor design and start to explore the vast design space of on-chip processing. The authors cover a number of different types of processor core.

As computing devices proliferate, demand increases for an understanding of emerging computing paradigms and models based on natural phenomena. Neural networks, evolution-based models, quantum computing, and DNA-based computing and simulations are all a necessary part of modern computing analysis and systems development. Vast literature exists on these new paradigms and their implications for a wide array of applications. This comprehensive handbook, the first of its kind to address the connection between nature-inspired and traditional computational paradigms, is a repository of case studies dealing with different problems in computing and solutions to these problems based on nature-inspired paradigms. The "Handbook of Nature-Inspired and Innovative Computing: Integrating Classical Models with Emerging Technologies" is an essential compilation of models, methods, and algorithms for researchers, professionals, and advanced-level students working in all areas of computer science, IT, biocomputing, and network engineering.

The current embedded processors often do not satisfy increasingly demanding computation requirements of embedded applications within acceptable energy efficiency, whereas application-specific integrated circuits require excessive design costs. In the Stanford Elm project, it was identified that instruction and data delivery, not computation, dominate the energy consumption of embedded processors. Consequently, the energy efficiency of delivering instructions and data must be sufficiently improved to close the efficiency gap between application-specific integrated circuits and programmable embedded processors. This dissertation demonstrates that the compiler and run-time system can play a crucial role in improving the energy efficiency of delivering instructions and data. Regarding instruction delivery, I present a compiler algorithm that manages L0 instruction scratch-pad memories that reside between processor cores and L1 caches. Despite the lack of tags, the scratch-pad memories with our algorithm can achieve lower miss rates than caches with the same capacities, saving significant instruction delivery energy. Regarding data delivery, I present methods that minimize memory-space requirements for parallelizing stream applications, applications that are commonly found in the embedded domain. When stream applications are parallelized in pipelining, large enough buffers are required between pipeline stages to sustain the throughput (e.g., double buffering). For static stream applications where production and consumption rates of stages are close to compile-time constants, a compiler analysis is presented, which computes the minimum buffer capacity that maximizes the throughput. Based on this analysis, a new static streamscheduling algorithm is developed, which yields considerable speed-up and data delivery energy saving compared to a previous algorithm. For dynamic stream applications, I present a dynamically-sized array-based queue design that achieves speed-up and data delivery energy saving compared to a linked-list based queue design.

The number of gates on a chip is quickly growing toward and beyond the one billion mark. Keeping all the gates running at the beat of a single or a few rationally related clocks is becoming impossible. In static timing analysis process variations and signal integrity issues stretch the timing margins to the point where they become too conservative and result in significant overdesign. Importance and difficulty of such problems push some developers to once again turn to asynchronous alternatives. However, the electronics industry for the most part is still reluctant to adopt asynchronous design (with a few notable exceptions) due to a common belief that we still lack a commercial-quality Electronic Design Automation tools (similar to the synchronous RTL-to-GDSII flow) for asynchronous circuits. The purpose of this paper is to counteract this view by presenting design flows that can tackle large designs without significant changes with respect to synchronous design flow. We are limiting ourselves to four design flows that we believe to be closest to this goal. We start from the Tangram flow, because it is the most commercially proven and it is one of the oldest from a methodological point of view. The other three flows (Null Convention Logic, de-synchronization, and gate-level pipelining) could be considered together as asynchronous re-implementations of synchronous (RTL- or gate-level) specifications. The main common idea is substituting the global clocks by local synchronizations. Their most important aspect is to open the possibility to implement large legacy synchronous designs in an almost "push button" manner, where all asynchronous machinery is hidden, so that synchronous RTL designers do not need to be re-educated. These three flows offer a trade-off from very low overhead, almost synchronous implementations, to very high performance, extremely robust dual-rail pipelines.

This book carefully details design tools and techniques for high-performance ASIC design. Using these techniques, the performance of ASIC designs can be improved by two to three times. Important topics include: Improving performance through microarchitecture; Timing-driven floorplanning; Controlling and exploiting clock skew; High performance latch-based design in an ASIC methodology; Automatically identifying and synthesizing complex logic gates; Automated cell sizing to increase performance and reduce power; Controlling process variation.These techniques are illustrated by designs running two to three times the speed of typical ASICs in the same process generation.

FPGA Architecture: Survey and Challenges reviews the historical development of programmable logic devices, the fundamental programming technologies that the programmability is built on, and then describes the basic understandings gleaned from research on architectures. It is an invaluable reference for engineers and computer scientists. It is also an excellent primer for senior or graduate-level students in electrical engineering or computer science.

Implementing energy-efficient CPUs and peripherals as well as reducing resource consumption have become emerging trends in computing. As computers increase in speed and power, their energy issues become more and more prevalent. The need to develop and promote environmentally friendly computer technologies and systems has also come to the forefront

The fact that there are more embedded computers than general-purpose computers and that we are impacted by hundreds of them every day is no longer news. What is news is that their increasing performance requirements, complexity and capabilities demand a new approach to their design. Fisher, Faraboschi, and Young describe a new age of embedded computing design, in which the processor is central, making the approach radically distinct from contemporary practices of embedded systems design. They demonstrate why it is essential to take a computing-centric and system-design approach to the traditional elements of nonprogrammable components, peripherals, interconnects and buses. These elements must be unified in a system design with high-performance processor architectures, microarchitectures and compilers, and with the compilation tools, debuggers and simulators needed for application development. In this landmark text, the authors apply their expertise in highly interdisciplinary hardware/software development and VLIW processors to illustrate this change in embedded computing. VLIW architectures have long been a popular choice in embedded systems design, and while VLIW is a running theme throughout the book, embedded computing is the core topic. Embedded Computing examines both in a book filled with fact and opinion based on the authors many years of R&D experience. · Complemented by a unique, professional-quality embedded tool-chain on the authors' website, http://www.vliw.org/book · Combines technical depth with real-world experience · Comprehensively explains the differences between general purpose computing systems and embedded systems at the hardware, software, tools and operating system levels. · Uses concrete examples to explain and motivate the trade-offs.

This book constitutes the thoroughly refereed post-conference proceedings of 19th International Workshop on Power and Timing Modeling, Optimization and Simulation, PATMOS 2009, featuring Integrated Circuit and System Design, held in Delft, The Netherlands during September 9-11, 2009. The 26 revised full papers and 10 revised poster papers presented were carefully reviewed and selected from numerous submissions. The papers are organized in topical sections on variability & statistical timing, circuit level techniques, power management, low power circuits & technology, system level techniques, power & timing optimization techniques, self-timed circuits, low power circuit analysis & optimization, and low power design studies.

Efficient design of embedded processors plays a critical role in embedded systems design. Processor description languages and their associated specification, exploration and rapid prototyping methodologies are used to find the best possible design for a given set of applications under various design constraints, such as area, power and performance. This book is the first, comprehensive survey of modern architecture description languages and will be an invaluable reference for embedded system architects, designers, developers, and validation engineers. Readers will see that the use of particular architecture description languages will lead to productivity gains in designing particular (application-specific) types of embedded processors. * Comprehensive coverage of all modern architecture description languages... use the right ADL to design your processor to fit your application; * Most up-to-date information available about each architecture description language from the developers...save time chasing down reliable documentation; * Describes how each architecture desccription language enables key design automation tasks, such as simulation, synthesis and testing...fit the ADL to your design cycle;

The second of two volumes in the Electronic Design Automation for Integrated Circuits Handbook, Second Edition, Electronic Design Automation for IC Implementation, Circuit Design, and Process Technology thoroughly examines real-time logic (RTL) to GDSII (a file format used to transfer data of semiconductor physical layout) design flow, analog/mixed signal design, physical verification, and technology computer-aided design (TCAD). Chapters contributed by leading experts authoritatively discuss design for manufacturability (DFM) at the nanoscale, power supply network design and analysis, design modeling, and much more. New to This Edition: Major updates appearing in the initial phases of the design flow, where the level of abstraction keeps rising to support more functionality with lower non-recurring engineering (NRE) costs Significant revisions reflected in the final phases of the design flow, where the complexity due to smaller and smaller geometries is compounded by the slow progress of shorter wavelength lithography New coverage of cutting-edge applications and approaches realized in the decade since publication of the previous edition—these are illustrated by new chapters on 3D circuit integration and clock design Offering improved depth and modernity, Electronic Design Automation for IC Implementation, Circuit Design, and Process Technology provides a valuable, state-of-the-art reference for electronic design automation (EDA) students, researchers, and professionals.

"This book has collected the latest research within the field of real-time systems engineering, and will serve as a vital reference compendium for practitioners and academics"--Provided by publisher.

This volume contains the proceedings of the Third International Conference on Trust and Trustworthy Computing (TRUST), held at the Ritz-Carlton hotel in Berlin, Germany, June 21–23, 2010. TRUST is a rapidly growing forum for research on the technical and soc- economic aspects of trustworthy infrastructures. TRUST provides an interdis- plinary forum for researchers, practitioners, and decision makers to explore new ideas and discuss experiences in building, designing, using, and understanding trustworthy computing systems. The third edition of TRUST welcomed manuscripts in two di?erent tracks: a Technical Strand and a Socio-economic Strand. We assembled an engaging program with 21 peer-reviewed technical papers and nine peer-reviewed soc- economic papers; eight keynotes from industry, academia, and government; and panel discussions on privacy and standards. In addition, this year, TRUST was co-located with four workshops: Trust in Cloud, Hardware Security, Emerging and Future Risks, and Anonymous Signatures. We would like to thank numerous individuals for their e?ort and contri- tion to the conference and for making TRUST 2010 possible: the Organizing Committee members—Nadine Palacios and Marcel Winandy—for their trem- dous help with all aspects of the organization;the Technicaland Socio-economic Program Committee members, whose names are listed on the following pages, together with the names of external reviewers who helped us in the process of selecting manuscripts to be included in the conference proceedings; the keynote and invited speakers; and the invited panel speakers.

Single-threaded software applications have ceased to see signi?cant gains in p- formance on a general-purpose CPU, even with further scaling in very large scale integration (VLSI) technology. This is a signi?cant problem for electronic design automation (EDA) applications, since the design complexity of VLSI integrated circuits (ICs) is continuously growing. In this research monograph, we evaluate custom ICs, ?eld-programmable gate arrays (FPGAs), and graphics processors as platforms for accelerating EDA algorithms, instead of the general-purpose sing- threaded CPU. We study applications which are used in key time-consuming steps of the VLSI design ?ow. Further, these applications also have different degrees of inherent parallelism in them. We study both control-dominated EDA applications and control plus data parallel EDA applications. We accelerate these applications on these different hardware platforms. We also present an automated approach for accelerating certain uniprocessor applications on a graphics processor. This monograph compares custom ICs, FPGAs, and graphics processing units (GPUs) as potential platforms to accelerate EDA algorithms. It also provides details of the programming model used for interfacing with the GPUs.

Field-programmable gate arrays (FPGAs), which are pre-fabricated, programmable digital integrated circuits (ICs), provide easy access to state-of-the-art integrated circuit process technology, and in doing so, democratize this technology of our time. This book is about comparing the qualities of FPGA – their speed performance, area and power consumption, against custom-fabricated ICs, and exploring ways of mitigating their de ciencies. This work began as a question that many have asked, and few had the resources to answer – how much worse is an FPGA compared to a custom-designed chip? As we dealt with that question, we found that it was far more dif cult to answer than we anticipated, but that the results were rich basic insights on fundamental understandings of FPGA architecture. It also encouraged us to nd ways to leverage those insights to seek ways to make FPGA technology better, which is what the second half of the book is about. While the question “How much worse is an FPGA than an ASIC?” has been a constant sub-theme of all research on FPGAs, it was posed most directly, some time around May 2004, by Professor Abbas El Gamal from Stanford University to us – he was working on a 3D FPGA, and was wondering if any real measurements had been made in this kind of comparison. Shortly thereafter we took it up and tried to answer in a serious way.

The book deals with engineering aspects of the two emerging and intertwined fields of synthetic and systems biology. Both fields hold promise to revolutionize the way molecular biology research is done, the

way today's drug discovery works and the way bio-engineering is done. Both fields stress the importance of building and characterizing small bio-molecular networks in order to synthesize incrementally and understand large complex networks inside living cells. Reminiscent of computer-aided design (CAD) of electronic circuits, abstraction is believed to be the key concept to achieve this goal. It allows hiding the overwhelming complexity of cellular processes by encapsulating network parts into abstract modules. This book provides a unique perspective on how concepts and methods from CAD of electronic circuits can be leveraged to overcome complexity barrier perceived in synthetic and systems biology.