经 典 原 版 书 库

复杂SoC设计

(英文版)

Engineering the Complex SOC

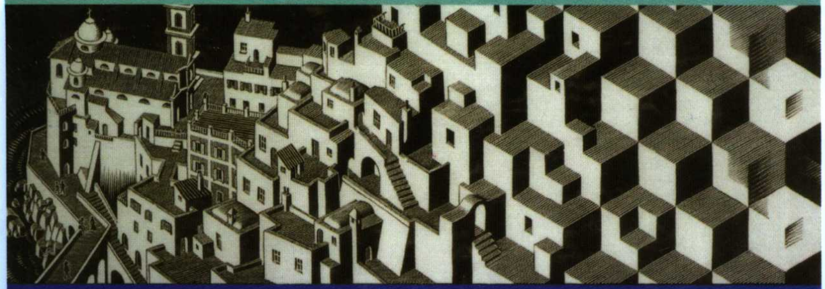

Fast, Flexible Design with Configurable Processors

CHRIS ROWEN

(美) Chris Rowen 著

机械工业出版社
China Machine Press

经 典 原 版 书 库

复杂SoC设计

（英文版）

Engineering the Complex SOC
Fast, Flexible Design with Configurable Processors

（美） Chris Rowen 著

机 械 工 业 出 版 社
China Machine Press

本书版权登记号：图字：01-2005-3622

图书在版编目（CIP）数据

复杂SoC设计（英文版）／（美）罗恩（Rowen, C.）著. －北京：机械工业出版社，2005.9

（经典原版书库）

书名原文：Engineering the Complex SOC: Fast, Flexible Design with Configurable Processors

ISBN 7-111-17193-4

Ⅰ. 复…　Ⅱ. 罗…　Ⅲ. 微处理器－系统设计－英文　Ⅳ. TP332

中国版本图书馆CIP数据核字（2005）第093805号

机械工业出版社（北京市西城区百万庄大街22号　邮政编码　100037）

责任编辑：迟振春

北京牛山世兴印刷厂印刷 · 新华书店北京发行所发行

2005年9月第1版第1次印刷

787mm × 1092mm 1/16 · 30.25印张

印数：0 001 -3 000册

定价：55.00元

出版者的话

文艺复兴以降，源远流长的科学精神和逐步形成的学术规范，使西方国家在自然科学的各个领域取得了垄断性的优势；也正是这样的传统，使美国在信息技术发展的六十多年间名家辈出、独领风骚。在商业化的进程中，美国的产业界与教育界越来越紧密地结合，计算机学科中的许多泰山北斗同时身处科研和教学的最前线，由此而产生的经典科学著作，不仅擘划了研究的范畴，还揭橥了学术的源变，既遵循学术规范，又自有学者个性，其价值并不会因年月的流逝而减退。

近年，在全球信息化大潮的推动下，我国的计算机产业发展迅猛，对专业人才的需求日益迫切。这对计算机教育界和出版界都既是机遇，也是挑战；而专业教材的建设在教育战略上显得举足轻重。在我国信息技术发展时间较短、从业人员较少的现状下，美国等发达国家在其计算机科学发展的几十年间积淀的经典教材仍有许多值得借鉴之处。因此，引进一批国外优秀计算机教材将对我国计算机教育事业的发展起积极的推动作用，也是与世界接轨、建设真正的世界一流大学的必由之路。

机械工业出版社华章图文信息有限公司较早意识到“出版要为教育服务”。自1998年开始，华章公司就将工作重点放在了遴选、移译国外优秀教材上。经过几年的不懈努力，我们与Prentice Hall，Addison-Wesley，McGraw-Hill，Morgan Kaufmann等世界著名出版公司建立了良好的合作关系，从它们现有的数百种教材中甄选出Tanenbaum，Stroustrup，Kernighan，Jim Gray等大师名家的一批经典作品，以“计算机科学丛书”为总称出版，供读者学习、研究及庋藏。大理石纹理的封面，也正体现了这套丛书的品位和格调。

“计算机科学丛书”的出版工作得到了国内外学者的鼎力襄助，国内的专家不仅提供了中肯的选题指导，还不辞劳苦地担任了翻译和审校的工作；而原书的作者也相当关注其作品在中国的传播，有的还专程为其书的中译本作序。迄今，“计算机科学丛书”已经出版了近百个品种，这些书籍在读者中树立了良好的口碑，并被许多高校采用为正式教材和参考书籍，为进一步推广与发展打下了坚实的基础。

随着学科建设的初步完善和教材改革的逐渐深化，教育界对国外计算机教材的需求和应用都步入一个新的阶段。为此，华章公司将加大引进教材的力度，在“华章教育”的总规划之下出版三个系列的计算机教材：除“计算机科学丛书”之外，对影印版的教材，则单独开辟出“经典原版书库”；同时，引进全美通行的教学辅导书“Schaum's Outlines”系列组成“全美经典学习指导系列”。为了保证这三套丛书的权威性，同时也为了更好地为学校和老师们服务，华章公司聘请了中国科学院、北京大学、清华大学、国防科技大学、复旦大学、上海交通大学、南京大学、浙江大学、中国科技大学、哈尔滨工业大学、西安交通大学、中国人民大学、北京航空航天大学、北京邮电大学、中山大学、解放军理工大学、郑州大学、湖北工学院、中国国家信息安全测评认证中心等国内重点大学和科研机构在计算机的各个领域的著名学者组成“专

家指导委员会”，为我们提供选题意见和出版监督。

这三套丛书是响应教育部提出的使用外版教材的号召，为国内高校的计算机及相关专业的教学度身订造的。其中许多教材均已为M. I. T.，Stanford，U.C. Berkeley，C. M. U. 等世界名牌大学所采用。不仅涵盖了程序设计、数据结构、操作系统、计算机体系结构、数据库、编译原理、软件工程、图形学、通信与网络、离散数学等国内大学计算机专业普遍开设的核心课程，而且各具特色——有的出自语言设计者之手、有的历经三十年而不衰、有的已被全世界的几百所高校采用。在这些圆熟通博的名师大作的指引之下，读者必将在计算机科学的宫殿中由登堂而入室。

权威的作者、经典的教材、一流的译者、严格的审校、精细的编辑，这些因素使我们的图书有了质量的保证，但我们的目标是尽善尽美，而反馈的意见正是我们达到这一终极目标的重要帮助。教材的出版只是我们的后续服务的起点。华章公司欢迎老师和读者对我们的工作提出建议或给予指正，我们的联系方法如下：

电子邮件：hzjsj@hzbook.com
联系电话：（010）68995264
联系地址：北京市西城区百万庄南街1号
邮政编码：100037

专家指导委员会

（按姓氏笔画顺序）

尤晋元	王　珊	冯博琴	史忠植	史美林
石教英	吕　建	孙玉芳	吴世忠	吴时霖
张立昂	李伟琴	李师贤	李建中	杨冬青
邵维忠	陆丽娜	陆鑫达	陈向群	周伯生
周克定	周傲英	孟小峰	岳丽华	范　明
郑国梁	施伯乐	钟玉琢	唐世渭	袁崇义
高传善	梅　宏	程　旭	程时端	谢希仁
裘宗燕	戴　葵			

List of Figures

FOREWORD

This is an important and useful book – important because it addresses a phenomenon that affects every industry sooner or later, and useful because it offers a clear, step-by-step methodology by which engineers and executives in the microelectronics industry can create growth and profit from this phenomenon. The companies that seize upon this opportunity will transform the way competition occurs in this industry. I wish that a hands-on guide such as this were available to strategists and design engineers in other industries where this phenomenon is occurring – industries as diverse as operating system software, automobiles, telecommunications equipment, and management education of the sort that we provide at the Harvard Business School. This industry-transforming phenomenon is called a change in the basis of competition – a change in the sorts of improvements in products and services that customers will willingly pay higher prices to get.

There is a natural and predictable process by which this change affects an industry. Chris Rowen, a gifted strategist, engineer and entrepreneur, has worked with me for several years to understand this process. It begins when a company develops a proprietary product that, while not good enough, comes closer to satisfying customers' needs than any of its competitors. The most successful firms do this through a proprietary and optimized architecture, because at this stage the functionality and reliability of such products are superior to those that employ an open, modular architecture. In order to provide proprietary, architecturally interdependent products, the most successful companies must be vertically integrated.

As the company strives to keep ahead of its direct competitors, however, it eventually overshoots the functionality and reliability that customers in less-demanding tiers of the market can utilize. This precipitates a change in the basis of competition in those tiers. Customers will no longer pay premium prices for better, faster and more reliable products, because they can't use those improvements. What is not good enough then becomes speed and convenience. Customers begin demanding new products that are responsively custom-configured to their needs, designed and delivered as rapidly and conveniently as possible. Innovations on these new trajectories of improvement are the improvements that merit attractive prices and drive changes in market share. In order to compete in this way – to be fast, flexible and responsive, the dominant architecture of the products must evolve toward a modular architecture, whose components and sub-systems interface according to industry standards. When this happens, there is no advantage to being integrated. Suppliers of components and sub-systems can begin developing, making and selling their products independently, dealing with partners and customers at arms' length because the key interface standards are completely and clearly specified. This condition begins at the bottom of the market, where functional overshoot occurs first, and then moves up inexorably to affect the higher tiers.

When the architecture of a product or a sub-system is modular, it is conformable. This conformability is very important when it is being incorporated within a next-level product whose functionality and reliability aren't yet good enough to fully satisfy what customers need. Modular conformability enables the customer to customize what it buys, getting every piece of functionality it needs, and none of the functionality it doesn't need.

The microprocessor industry is going through precisely this transition. Microprocessors, and the size of the features from which they are built, historically have not been good enough – and as a result, their architectures have been proprietary and optimized. Now, however, there is strong evidence that for mainstream tiers of the market the basis of competition is changing. Microprocessors are more than fast enough for most computer users. In pursuit of Moore's Law, circuit fabricators have shrunk feature sizes to such a degree that in most tiers of the market, circuit designers are awash in transistors. They cannot use all the transistors that Moore's Law has made available. As a result, especially in embedded, mobile and wireless applications, custom-configured processors are taking over. Their modular configurability helps designers to optimize the performance of the product systems in which the processors are embedded.

With this change, the pace of the microprocessor industry is accelerating. Product design cycles, which in the era of interdependent architectures had to be measured in years, are collapsing to months. In the future they will be counted in weeks. Clean, modular interfaces between the modules that comprise a circuit – libraries of reusable IP – ultimately will enable engineers who are not experts in processor design, to build their own circuits. Ultimately software engineers will be able to design processors that are custom-configured to optimize the performance of their software application.

Clearer interface standards are being defined between circuit designers and fabs, enabling a dis-integrated industry structure to overtake the original integrated one. This structure began at the low end, and now dominates the mainstream of the market as well. The emergence of the modular "designed-to-order" processor is an important milestone in the evolution of the microprocessor industry, but modular processors have broader implications. Just as the emergence of the personal computer allowed a wide range of workers to computerize their tasksfor the first time, the configurable processor will change the lives of a wide range of chip designers and chip users. This new processor-based methodology enables these designers and users to specify and program processors for tasks that are too sensitive to cost or energy-efficiency for traditional processors. It empowers the ordinary software or hardware developer to create new computing engines, once the province of highly-specialized microprocessor architecture and development teams. And these new processor blocks are likely to be used in large numbers – tens and hundreds per chip, with total configurable processors vastly outnumbering traditional microprocessors. The future will therefore be very different from the past.

The design principles and techniques that Chris Rowen describes in this book will be extremely useful to companies that want to capitalize on these changes to create new growth. I thank him for this gift to the semiconductor industry.

Clayton M. Christensen
Robert and Jane Cizik Professor of Business Administration
Harvard Business School
March 2004

FOREWORD

For more than 30 years, Moore's law has been a driving force in the computing and electronics industries, constantly forcing changes and driving innovation as it provides the means to integrate ever-larger systems onto a single chip. Moore's law was the driving force that led to the microprocessor's dominance throughout the world of computing, beginning about 20 years ago. Today, Moore's law is the driver behind the system-on-chip (SOC) paradigm that uses the vastly increased transistor density to integrate ever-larger system components onto a single chip, thereby reducing cost, power consumption, and physical size.

Rowen structures this thorough treatise on the design of complex SOCs around six fundamental problems. These range from market forces, such as tight time to market and limited volume before obsolescence, to the technical challenges of achieving acceptable performance and cost while adhering to an aggressive schedule. These six challenges lead to a focus on specific parts of the SOC design process, from the integration of application-specific logic to the maximization of performance/cost ratios through processor customization.

As Rowen clearly illustrates, the processor and its design is at the core of any complex SOC. After all, a software-mostly solution is likely to reduce implementation time and risk; the difficulty is that such solutions are often unable to achieve acceptable performance or efficiency. For most applications, some combination of a processor (executing application-specific code) and application-specific hardware are necessary. In Chapter 4, the author deals with the critical issue of interfacing custom hardware and embedded processor cores by observing that a mix of custom communication and interconnection hardware together with software to handle decision making and less common tasks often enables the designer to achieve the desired balance of implementation effort, risk, performance, and cost.

Chapters 5 and 6 form the core of this book and build on the years of experience that Tensilica has had in building SOCs based on customized processors. Although the potential advantages of designing a customized processor should be clear—lower cost and better performance—the required CAD tools, customizable building blocks, and software were previously unavailable. Now they are, and Rowen discusses how to undertake the design of both the software and hardware for complex SOCs using state-of-the art tools.

Chapters 7 and 8 address the challenge of obtaining high performance in such systems. For processors, performance comes primarily from the exploitation of parallelism. This is a central topic in both these chapters. Chapter 7 discusses the use of pipelining to achieve higher performance within a single instruction flow in a single processor. Chapter 8 looks to the future, which will increasingly make use of multiple processors, configured and connected according to the needs of the application. The use of multiple processors broadly represents the future in high-performance computer architecture, and not just in embedded applications. I am delighted to see that SOCs are playing a key role in charting this future.

The design of SOCs for new and challenging applications—ranging from telecommunications, to information appliances, to applications we have yet to dream of—is creating new opportunities for computing. This well-written and comprehensive book will help you be a successful participant in these exciting endeavors.

John Hennessy
President
Stanford University
March 2004