# Networks On Chips Technology And Tools Systems On Silicon

This book constitutes the refereed proceedings of the 24th International Conference on Architecture of Computing Systems, ARCS 2011, held in Lake Como, Italy, in February 2011. The 22 revised full papers presented in seven technical sessions were carefully reviewed and selected from 62 submissions. The papers are organized in topical sections on customization and application specific accelerators; multi/many-core architectures; adaptive system architectures; processor architectures; memory architectures optimization; organic and autonomic computing; network-on-chip architectures.

The early era of neural network hardware design (starting at 1985) was mainly technology driven. Designers used almost exclusively analog signal processing concepts for the recall mode. Learning was deemed not to cause a problem because the number of implementable synapses was still so low that the determination of weights and thresholds could be left to conventional computers. Instead, designers tried to directly map neural parallelity into hardware. The architectural concepts were accordingly simple and produced the so called interconnection problem which, in turn, made many engineers believe it could be solved by optical implementation in adequate fashion only. Furthermore, the inherent fault-tolerance and limited computation accuracy of neural networks were claimed to justify that little effort is to be spend on careful design, but most effort be put on technology issues. As a result, it was almost impossible to predict whether an electronic neural network would function in the way it was simulated to do. This limited the use of the first neuro-chips for further experimentation, not to mention that real-world applications called for much more synapses than could be implemented on a single chip at that time. Meanwhile matters have matured. It is recognized that isolated definition of the effort of analog multiplication, for instance, would be just as inappropriate on the part ofthe chip designer as determination of the weights by simulation, without allowing for the computing accuracy that can be achieved, on the part of the user.

In 2007 The Design, Automation and Test in Europe (DATE) conference celebrated its tenth anniversary. As a tribute to the chip and system-level design and design technology community, this book presents a compilation of the three most influential papers of each year. This provides an excellent historical overview of the evolution of a domain that contributed substantially to the growth and competitiveness of the circuit electronics and systems industry.

This book provides comprehensive coverage of the dependability challenges in today's advanced computing systems. It is an in-depth discussion of all the technological and design-level techniques that may be used to overcome these issues and analyzes various dependability-assessment methods. The impact of individual application scenarios on the definition of challenges and solutions is considered so that the designer can clearly assess the problems and adjust the solution based on the specifications in question. The book is composed of three sections, beginning with an introduction to current dependability challenges arising in complex computing systems implemented with nanoscale technologies, and of the effect of the application scenario. The second section details all the fault-tolerance techniques that are applicable in the manufacture of reliable advanced computing devices. Different levels, from technology-level fault avoidance to the use of error correcting codes and system-level checkpointing are introduced and explained as applicable to the different application scenario requirements. Finally the third section proposes a roadmap of future trends in and perspectives on the dependability and manufacturability of advanced computing systems from the special point of view of industrial stakeholders. Dependable Multicore Architectures at Nanoscale showcases the original ideas and concepts introduced into the field of nanoscale manufacturing and systems reliability over nearly four years of work within COST Action

IC1103 MEDIAN, a think-tank with participants from 27 countries. Academic researchers and graduate students working in multi-core computer systems and their manufacture will find this book of interest as will industrial design and manufacturing engineers working in VLSI companies.

The first book to survey this emerging field in digital system design.

A professional reference that examines the gigabit per second computer networks that make it possible to share vast quantities of data among many computer systems. Key technologies, important protocols and applications, and the practical issues involved in implementing gigabit networks are all addressed, and where research is still incomplete, important unsolved issues are presented. Could also be used as a textbook for a graduate course on gigabit networking. Annotation copyright by Book News, Inc., Portland, OR

?????:????

A presentation of state-of-the-art approaches from an industrial applications perspective, Communication Architectures for Systems-on-Chip shows professionals, researchers, and students how to attack the problem of data communication in the manufacture of SoC architectures. With its lucid illustration of current trends and research improving the performance, quality, and reliability of transactions, this is an essential reference for anyone dealing with communication mechanisms for embedded systems, systems-on-chip, and multiprocessor architectures—or trying to overcome existing limitations. Exploring architectures currently implemented in manufactured SoCs—and those being proposed—this book analyzes a wide range of applications, including: Well-established communication buses Less common networks-on-chip Modern technologies that include the use of carbon nanotubes (CNTs) Optical links used to speed up data transfer and boost both security and quality of service (QoS) The book’s contributors pay special attention to newer problems, including how to protect transactions of critical on-chip information (personal data, security keys, etc.) from an external attack. They examine mechanisms, revise communication protocols involved, and analyze overall impact on system performance.

Transactions on HiPEAC aims at the timely dissemination of research contributions in computer architecture and compilation methods for high-performance embedded computer systems. Recognizing the convergence of embedded and general-purpose computer systems, this journal publishes original research on systems targeted at specific computing tasks as well as systems with broad application bases. The scope of the journal therefore covers all aspects of computer architecture, code generation and compiler optimization methods of interest to researchers and practitioners designing future embedded systems. This 4th issue contains 21 papers carefully reviewed and selected out of numerous submissions and is divided in four sections. The first section contains five regular papers. The second section consists of the top four papers from the 4th International Conference on High-Performance Embedded Architectures and Compilers, HiPEAC 2009, held in Paphos, Cyprus, in January 2009. The third section contains a set of six papers providing a snap-shot from the Workshop on Software and Hardware Challenges of Manycore Platforms, SHCMP 2008 held in Beijing, China, in June 2008. The fourth section consists of six papers from the 8th IEEE International Symposium on Systems, Architectures, Modeling and Simulation, SAMOS VIII (2008) held in Samos, Greece, in July 2008.

Datacenter networks provide the communication substrate for large parallel

computer systems that form the ecosystem for high performance computing (HPC) systems and modern Internet applications. The design of new datacenter networks is motivated by an array of applications ranging from communication intensive climatology, complex material simulations and molecular dynamics to such Internet applications as Web search, language translation, collaborative Internet applications, streaming video and voice-over-IP. For both Supercomputing and Cloud Computing the network enables distributed applications to communicate and interoperate in an orchestrated and efficient way. This book describes the design and engineering tradeoffs of datacenter networks. It describes interconnection networks from topology and network architecture to routing algorithms, and presents opportunities for taking advantage of the emerging technology trends that are influencing router microarchitecture. With the emergence of "many-core" processor chips, it is evident that we will also need "many-port" routing chips to provide a bandwidth-rich network to avoid the performance limiting effects of Amdahl's Law. We provide an overview of conventional topologies and their routing algorithms and show how technology, signaling rates and cost-effective optics are motivating new network topologies that scale up to millions of hosts. The book also provides detailed case studies of two high performance parallel computer systems and their networks. Table of Contents: Introduction / Background / Topology Basics / High-Radix Topologies / Routing / Scalable Switch Microarchitecture / System Packaging / Case Studies / Closing Remarks

This book contains extended and revised versions of the best papers presented at the 17th IFIP WG 10.5/IEEE International Conference on Very Large Scale Integration, VLSI-SoC 2009, held in Florianópolis, Brazil, in October 2009. The 8 papers included in the book together with two keynote talks were carefully reviewed and selected from 27 papers presented at the conference. The papers cover a wide variety of excellence in VLSI technology and advanced research addressing the current trend toward increasing chip integration and technology process advancements bringing about stimulating new challenges both at the physical and system-design levels, as well as in the test of theses systems.

Developing NoC based interconnect tailored to a particular application domain, satisfying the application performance constraints with minimum power-area overhead is a major challenge. With technology scaling, as the geometries of on-chip devices reach the physical limits of operation, another important design challenge for NoCs will be to provide dynamic (run-time) support against permanent and intermittent faults that can occur in the system. The purpose of Designing Reliable and Efficient Networks on Chips is to provide state-of-the-art methods to solve some of the most important and time-intensive problems encountered during NoC design.

2nd Edition of Quantitative Finance and Risk Management: A Physicist's ApproachWritten by a physicist with over 15 years of experience as a quant on Wall Street, this book treats a wide variety of topics. Presenting the theory and

practice of quantitative finance and risk, it delves into the “how to” and “what it's like” aspects not covered in textbooks or research papers. Both standard and new results are presented. A “Technical Index” indicates the mathematical level — from zero to PhD — for each chapter. The finance in each chapter is self-contained. Real-life comments on “life as a quant” are included.An errata and Additions (3rd Reprint, 2008) to the book is available.

Networks-on-Chip: From Implementations to Programming Paradigms provides a thorough and bottom-up exploration of the whole NoC design space in a coherent and uniform fashion, from low-level router, buffer and topology implementations, to routing and flow control schemes, to co-optimizations of NoC and high-level programming paradigms. This textbook is intended for an advanced course on computer architecture, suitable for graduate students or senior undergrads who want to specialize in the area of computer architecture and Networks-on-Chip. It is also intended for practitioners in the industry in the area of microprocessor design, especially the many-core processor design with a network-on-chip. Graduates can learn many practical and theoretical lessons from this course, and also can be motivated to delve further into the ideas and designs proposed in this book. Industrial engineers can refer to this book to make practical tradeoffs as well. Graduates and engineers who focus on off-chip network design can also refer to this book to achieve deadlock-free routing algorithm designs. Provides thorough and insightful exploration of NoC design space. Description from low-level logic implementations to co-optimizations of high-level program paradigms and NoCs. The coherent and uniform format offers readers a clear, quick and efficient exploration of NoC design space Covers many novel and exciting research ideas, which encourage researchers to further delve into these topics. Presents both engineering and theoretical contributions. The detailed description of the router, buffer and topology implementations, comparisons and analysis are of high engineering value.

This book presents design techniques, analysis and implementation of high performance and power efficient, variation tolerant on-chip interconnects. Given the design paradigm shift to multi-core, interconnect-centric designs and the increase in sources of variability and their impact in sub-100nm technologies, this book will be an invaluable reference for anyone concerned with the design of next generation, high-performance electronics systems.

For more than 20 years, Network World has been the premier provider of information, intelligence and insight for network and IT executives responsible for the digital nervous systems of large organizations. Readers are responsible for designing, implementing and managing the voice, data and video systems their companies use to support everything from business critical applications to employee collaboration and electronic commerce.

In the network economy, concepts of management knowledge, management learning, and business school organization should change. Otherwise, they will not survive the 21st century. Different (f)actors are putting new demand upon providers of

management education and traditional providers of management education are faced with new competitors. Moreover, the dynamics of the playing field have changed, as have approaches to (management) learning. Management Education in the Network Economy proposes the idea of networked business school to cope with these challenges. The book deals with the following subjects: 1) Current economic and organizational realities can best be viewed from the perspective of network organization; management knowledge and education should reflect these transformations to survive. 2) The idea and organization of (management) learning are revolutionizing, as well as the market for (management) education, which brings about huge changes for business schools. 3) Business school, particularly, should capitalize on these transformations and should strategically (re)organize and (re)position themselves to compete in the playing field for management education. 4) A networked learning environment is an integrative and effective learning environment for organizing management education in the 21st century network economy. 5) The networked business school is the organizational form to survive in the 21st century network economy, reflecting the environmental changes and demands, and to realize a competitive edge in the field of management education.

The day when fiber will deliver new, yet now only foreseeable, broadband ser vices to the end user is getting nearer and nearer as we make our way towards the prophetic year 2000. Step by step, as we move from first generation lasers and fibers to the by now common erbium-doped fiber amplifiers, looking forward to such things as wavelength multiplexing and solitons, photonic switching and optical storage, the community of researchers in optical communications has stepped into the era of photonic networks. It is not just a question of terminology. Optical communication means tech nology to the same extent that photonic network means services. If it is true that information is just as marketable a product as oil or coke, the providing of an extensive global information infrastructure may end up having an even greater impact than the setting up of a world-wide railroad network did at the beginning of the industrial era. Just like wagons, bandwidth will be responsible for carrying and delivering goods to customers. The challenge for all of us in this field is for it to function in every section of the overall network, transport, access and customer area, in the best possible way: the fastest, most economical and most flexible. New services provided by a new network that exploits the potential and peculiarities of photonics surely requires a rethinking of solutions, new ideas, new architec tures, new design, especially where electronics is still dominant, as in transport and access networks.

The design of today's semiconductor chips for various applications, such as telecommunications, poses various challenges due to the complexity of these systems. These highly complex systems-on-chips demand new approaches to connect and manage the communication between on-chip processing and storage components and networks on chips (NoCs) provide a powerful solution. This book is the first to provide a unified overview of NoC technology. It includes in-depth analysis of all the on-chip communication challenges, from physical wiring implementation up to software architecture, and a complete classification of their various Network-on-Chip approaches and solutions. * Leading-edge research from world-renowned experts in academia and industry with state-of-the-art technology implementations/trends * An integrated presentation not currently available in any other book * A thorough introduction to

current design methodologies and chips designed with NoCs

Communication has become a bottleneck for modern microprocessors and multi-core chips because metal wires don't scale. The problem becomes worse as the number of components increases and chips become bigger. Traditional Systems-on-Chips (SoCs) interconnect architectures are based on shared-bus communication, which can carry only one communication transaction at a time. This limits the communication bandwidth and scalability. Networks-on-Chip (NoC) were proposed as a promising solution for designing large and complex SoCs. The NoC paradigm provides better scalability and reusability for future SoCs, however, long-distance multi-hop communication through traditional metal wires suffers from both high latency and power consumption. A radical solution to address this challenge is to add long-range, low power, and high-bandwidth single-hop links between distant cores. The use of optical or on-chip RF wireless links has been explored in this context. However, all previous work has focused on regular mesh-based metal wire fabrics that were expanded with one or two additional link types only for long-distance communication. In this thesis we address the following main research questions to address the above-mentioned challenges: (1) What library of different link types would represent an optimum in the design space? (2) How would these links be used to design an application-specific NoC architecture? (3) How would applications use the resulting NoC architecture efficiently? We hypothesize that networks with a higher degree of heterogeneity, i.e., three or more link types, will improve the network throughput and consume less energy compared to traditional NoC architectures. In order to verify our hypothesis and to address the research challenges, we design and analyze optimal heterogeneous networks under different realistic traffic models by considering different cost and performance trade-offs in a comprehensive technology-agnostic simulation framework that uses metaheuristic optimization techniques. As opposed to related work, our heterogeneous links can be placed anywhere in the network, which allows to explore the entire search space. The resulting application-specific networks are then analyzed by using complex network techniques, such as community detection and small-worldness, to understand how heterogeneous link types are used to improve the NoCs performance and cost. Next, we use the application-specific networks as a target architecture for other applications. The goal is to evaluate the performance of our new NoCs for applications they have not been designed for by finding optimal resource allocations. Our results show that there is an optimal number of heterogeneous link types for each set of constraints and that networks with three or more heterogeneous link types provide significantly higher throughput along with lower energy consumption compared to both homogeneous link type and regular 2D mesh networks under three different traffic scenarios. Our evolved networks with three different technology-driven link types, namely metal wires, wireless, and optical links, provide 15% more throughput and fourteen times less energy consumption compared to homogeneous link type network. When ten different abstract link types are used in the design, 12% more throughput and 52% less energy consumption are obtained compared to networks with three different technology-driven link types. This shows that heterogeneous NoC designs based on traditional metal wires, wireless, and optical links, occupy a non-optimal spot in the entire design space. Our results further show that heterogeneous NoCs scale up significantly better in terms of performance and cost compared to mesh networks. We uncovered that network

communities evolve robustly and that heterogeneous link types are efficiently establishing inter- and intra-subnet connections depending on their link type properties. We also show that mapping an application on our application-specific NoC architecture provides on average 45% more throughput at 70% less energy consumption compared to regular 2D mesh networks. The NoCs are therefore not only good for the application they were designed for, but for a broad range of other applications as well.

During the past few years there has been an dramatic upsurge in research and development, implementations of new technologies, and deployments of actual solutions and technologies in the diverse application areas of embedded systems. These areas include automotive electronics, industrial automated systems, and building automation and control. Comprising 48 chapters and the contributions of 74 leading experts from industry and academia, the Embedded Systems Handbook, Second Edition presents a comprehensive view of embedded systems: their design, verification, networking, and applications. The contributors, directly involved in the creation and evolution of the ideas and technologies presented, offer tutorials, research surveys, and technology overviews, exploring new developments, deployments, and trends. To accommodate the tremendous growth in the field, the handbook is now divided into two volumes. New in This Edition: Processors for embedded systems Processor-centric architecture description languages Networked embedded systems in the automotive and industrial automation fields Wireless embedded systems Embedded Systems Design and Verification Volume I of the handbook is divided into three sections. It begins with a brief introduction to embedded systems design and verification. The book then provides a comprehensive overview of embedded processors and various aspects of system-on-chip and FPGA, as well as solutions to design challenges. The final section explores power-aware embedded computing, design issues specific to secure embedded systems, and web services for embedded devices. Networked Embedded Systems Volume II focuses on selected application areas of networked embedded systems. It covers automotive field, industrial automation, building automation, and wireless sensor networks. This volume highlights implementations in fast-evolving areas which have not received proper coverage in other publications. Reflecting the unique functional requirements of different application areas, the contributors discuss inter-node communication aspects in the context of specific applications of networked embedded systems.

As the demand for digital communication networks has increased, so have the challenges in network component design. To meet ever-escalating performance, flexibility, and economy requirements, the networking industry has opted to build products around network processors. These new chips range from task-specific processors, such as classification and encryption engines, to more general-purpose packet or communications processors. Programmable yet application-specific, their designs are tailored to efficiently implement communications applications such as routing, protocol analysis, voice and data convergence, firewalls, VPNs, and QoS. Network processor design is an emerging field with issues and opportunities both numerous and formidable. To help meet this challenge, the editors of this volume created the first Workshop on Network Processors, a forum for scientists and engineers from academia and industry to discuss their latest research in the architecture, design, programming, and use of these devices. In addition to including the results of the

Workshop in this volume, the editors also present specially commissioned material from practicing designers, who discuss their companies' latest network processors. Network Processor Design: Issues and Practices is an essential reference on network processors for graduate students, researchers, and practicing designers. * Includes contributions from major academic and industrial research labs including Aachen University of Technology; Cisco Systems; Infineon Technologies; Intel Corp.; North Carolina State University; Swiss Federal Institute of Technology; University of California, Berkeley; University of Dortmund; University of Washington; and Washington University. * Examines the latest network processors from Agere Systems, Cisco, IBM, Intel, Motorola, Sierra Inc., and TranSwitch.

"The continuing advances in processing technology result in significant decreases in the feature size of integrated circuits. This aggressive transistor scaling enables integration of a large set of functionality inside a single Integrated Circuit (IC). As processing technology scales down, permanent and transient faults have become more frequent in new ICs. Consequently, along with opportunities to integrate a large number of processing elements, fault-tolerant architecture and enhanced debug infrastructure must be incorporated in new products. Network on Chips (NoCs) are poised to address the demands for high communication bandwidth among cores. A comprehensive study on fault-tolerant NoC routers and on-chip debug infrastructure are carried out in this thesis. This thesis presents fault-tolerant NoC router microarchitectures which can be can be incorporated in future hierarchical topology inside a chip. Armed with a new flow control mechanism, as well as an enhanced Virtual Channel (VC) regulator, the proposed router can mitigate the effect of both transient and permanent errors. The proposed router is the first study enabling inter-channel buffer sharing. In the realm of on-chip design instrumentation (post-silicon debugging), a tool that builds a synthesizable hierarchical trigger unit is presented. The proposed approach enables silicon debugging in a time-multiplexed fashion by producing several hierarchical trigger modules. These modules can be incorporated inside the limited silicon area. Compared to previous mechanisms, the detection of overlapped failure patterns can be carried out by 60-65 % reduction in hardware overhead. This thesis also proposes a new assertion-checker clustering algorithm along with several mechanisms to incorporate them into the on-chip debug infrastructure. The proposed debug infrastructure leads to better results in terms of the energy consumption and design coverage compared to previous work." --

This cutting-edge book on off-chip technologies puts the hottest breakthroughs in high-density compliant electrical interconnects, nanophotonics, and microfluidics at your fingertips, integrating the full range of mathematics, physics, and technology issues together in a single comprehensive source. You get full details on state-of-the-art I/O interconnects and packaging, including mechanically compliant I/O approaches, fabrication, and assembly, followed by the latest advances and applications in power delivery design, analysis, and modeling. The book explores interconnect structures, materials, and packages for achieving high-bandwidth off-chip electrical communication, including optical interconnects and chip-to-chip signaling approaches, and brings you up to speed on CMOS integrated optical devices, 3D integration, wafer stacking technology, and through-wafer interconnects.

This is a practical book for computer engineers who want to understand or implement

hardware/software systems. It focuses on problems that require one to combine hardware design with software design – such problems can be solved with hardware/software codesign. When used properly, hardware/software co- sign works better than hardware design or software design alone: it can improve the overall performance of digital systems, and it can shorten their design time. Hardware/software codesign can help a designer to make trade-offs between the ?exibility and the performanceof a digital system. To achieve this, a designer needs to combine two radically different ways of design: the sequential way of dec- position in time, using software, with the parallel way of decomposition in space, using hardware. Intended Audience This book assumes that you have a basic understandingof hardware that you are - miliar with standard digital hardware componentssuch as registers, logic gates, and components such as multiplexers and arithmetic operators. The book also assumes that you know how to write a program in C. These topics are usually covered in an introductory course on computer engineering or in a combination of courses on digital design and software engineering.

Over the past decade, system-on-chip (SoC) designs have evolved to address the ever increasing complexity of applications, fueled by the era of digital convergence. Improvements in process technology have effectively shrunk board-level components so they can be integrated on a single chip. New on-chip communication architectures have been designed to support all inter-component communication in a SoC design. These communication architecture fabrics have a critical impact on the power consumption, performance, cost and design cycle time of modern SoC designs. As application complexity strains the communication backbone of SoC designs, academic and industrial R&D efforts and dollars are increasingly focused on communication architecture design. On-Chip Communication Architecures is a comprehensive reference on concepts, research and trends in on-chip communication architecture design. It will provide readers with a comprehensive survey, not available elsewhere, of all current standards for on-chip communication architectures. A definitive guide to on-chip communication architectures, explaining key concepts, surveying research efforts and predicting future trends Detailed analysis of all popular standards for on-chip communication architectures Comprehensive survey of all research on communication architectures, covering a wide range of topics relevant to this area, spanning the past several years, and up to date with the most current research efforts Future trends that with have a significant impact on research and design of communication architectures over the next several years

Current multimedia and telecom applications require complex, heterogeneous multiprocessor system on chip (MPSoC) architectures with specific communication infrastructure in order to achieve the required performance. Heterogeneous MPSoC includes different types of processing units (DSP, microcontroller, ASIP) and different communication schemes (fast links, non standard memory organization and access). Programming an MPSoC requires the generation of efficient software running on MPSoC from a high level environment, by using the characteristics of the architecture. This task is known to be tedious and error prone, because it requires a combination of high level programming environments with low level software design. This book gives an overview of concepts related to embedded software design for MPSoC. It details a full software design approach, allowing systematic, high-level mapping of software applications on heterogeneous MPSoC. This approach is based on gradual refinement of hardware/software interfaces and simulation models allowing to validate the software at different abstraction levels. This book combines Simulink for high level programming and SystemC for the low level software development. This approach is illustrated with multiple examples of application software and MPSoC architectures that can be used for deep understanding of software design for MPSoC.

A Comprehensive, Thorough Introduction to High-Speed Networking Technologies and

Protocols Network Infrastructure and Architecture: Designing High-Availability Networks takes a unique approach to the subject by covering the ideas underlying networks, the architecture of the network elements, and the implementation of these elements in optical and VLSI technologies. Additionally, it focuses on areas not widely covered in existing books: physical transport and switching, the process and technique of building networking hardware, and new technologies being deployed in the marketplace, such as Metro Wave Division Multiplexing (MWDM), Resilient Packet Rings (RPR), Optical Ethernet, and more. Divided into five succinct parts, the book covers: Optical transmission Networking protocols VLSI chips Data switching Networking elements and design Complete with case studies, examples, and exercises throughout, the book is complemented with chapter goals, summaries, and lists of key points to aid readers in grasping the material presented. Network Infrastructure and Architecture offers professionals, advanced undergraduates, and graduate students a fresh view on high-speed networking from the physical layer perspective.

In Advanced ULSI interconnects – fundamentals and applications we bring a comprehensive description of copper-based interconnect technology for ultra-lar- scale integration (ULSI) technology for integrated circuit (IC) application. In- grated circuit technology is the base for all modern electronics systems. You can ?nd electronics systems today everywhere: from toys and home appliances to a- planes and space shuttles. Electronics systems form the hardware that together with software are the bases of the modern information society. The rapid growth and vast exploitation of modern electronics system create a strong demand for new and improved electronic circuits as demonstrated by the amazing progress in the ?eld of ULSI technology. This progress is well described by the famous “Moore’s law” which states, in its most general form, that all the metrics that describe integrated circuit performance (e. g. , speed, number of devices, chip area) improve expon- tially as a function of time. For example, the number of components per chip d- bles every 18 months and the critical dimension on a chip has shrunk by 50% every 2 years on average in the last 30 years. This rapid growth in integrated circuits te- nology results in highly complex integrated circuits with an increasing number of interconnects on chips and between the chip and its package. The complexity of the interconnect network on chips involves an increasing number of metal lines per interconnect level, more interconnect levels, and at the same time a reduction in the interconnect line critical dimensions.

Addresses the Challenges Associated with System-on-Chip Integration Network-on-Chip: The Next Generation of System-on-Chip Integration examines the current issues restricting chip-on-chip communication efficiency, and explores Network-on-chip (NoC), a promising alternative that equips designers with the capability to produce a scalable, reusable, and high-performance communication backbone by allowing for the integration of a large number of cores on a single system-on-chip (SoC). This book provides a basic overview of topics associated with NoC-based design: communication infrastructure design, communication methodology, evaluation framework, and mapping of applications onto NoC. It details the design and evaluation of different proposed NoC structures, low-power techniques, signal integrity and reliability issues, application mapping, testing, and future trends. Utilizing examples of chips that have been implemented in industry and academia, this text presents the full architectural design of components verified through implementation in industrial CAD tools. It describes NoC research and developments, incorporates theoretical proofs strengthening the analysis procedures, and includes algorithms used in NoC design and synthesis. In addition, it considers other upcoming NoC issues, such as low-power NoC design, signal integrity issues, NoC testing, reconfiguration, synthesis, and 3-D NoC design. This text comprises 12 chapters and covers: The evolution of NoC from SoC—its research and developmental challenges NoC protocols, elaborating flow control, available network topologies, routing mechanisms, fault tolerance, quality-of-service support, and the design of

network interfaces The router design strategies followed in NoCs The evaluation mechanism of NoC architectures The application mapping strategies followed in NoCs Low-power design techniques specifically followed in NoCs The signal integrity and reliability issues of NoC The details of NoC testing strategies reported so far The problem of synthesizing application-specific NoCs Reconfigurable NoC design issues Direction of future research and development in the field of NoC Network-on-Chip: The Next Generation of System-on-Chip Integration covers the basic topics, technology, and future trends relevant to NoC-based design, and can be used by engineers, students, and researchers and other industry professionals interested in computer architecture, embedded systems, and parallel/distributed systems.

The implementation of networks-on-chip (NoC) technology in VLSI integration presents a variety of unique challenges. To deal with specific design solutions and research hurdles related to intra-chip data exchange, engineers are challenged to invoke a wide range of disciplines and specializations while maintaining a focused approach. Leading Researchers Present Cutting-Edge Designs Tools Networks-on-Chips: Theory and Practice facilitates this process, detailing the NoC paradigm and its benefits in separating IP design and functionality from chip communication requirements and interfacing. It starts with an analysis of 3-D NoC architectures and progresses to a discussion of NoC resource allocation, processor traffic modeling, and formal verification, with an examination of protocols at different layers of abstraction. An exploration of design methodologies, CAD tool development, and system testing, as well as communication protocol, the text highlights important emerging research issues, such as Resource Allocation for Quality of Service (QoS) on-chip communication Testing, verification, and network design methodologies Architectures for interconnection, real-time monitoring, and security requirements Networks-on-Chip Protocols Presents a flexible MPSoC platform to easily implement multimedia applications and evaluate future video encoding standards This useful guide tackles power and energy issues in NoC-based designs, addressing the power constraints that currently limit the embedding of more processing elements on a single chip. It covers traffic modeling and discusses the details of traffic generators. Using unique case studies and examples, it covers theoretical and practical issues, guiding readers through every phase of system design.

Networks on ChipsTechnology and ToolsElsevier

This book contains extended and revised versions of the best papers that were p-sented during the 16th edition of the IFIP/IEEE WG10.5 International Conference on Very Large Scale Integration, a global System-on-a-Chip Design & CAD conference. The 16th conference was held at the Grand Hotel of Rhodes Island, Greece (October 13–15, 2008). Previous conferences have taken place in Edinburgh, Trondheim, V-couver, Munich, Grenoble, Tokyo, Gramado, Lisbon, Montpellier, Darmstadt, Perth, Nice and Atlanta. VLSI-SoC 2008 was the 16th in a series of international conferences sponsored by IFIP TC 10 Working Group 10.5 and IEEE CEDA that explores the state of the art and the new developments in the field of VLSI systems and their designs. The purpose of the conference was to provide a forum to exchange ideas and to present industrial and research results in the fields of VLSI/ULSI systems, embedded systems and - croelectronic design and test.

This book gives a broad look at both fundamental networking technology and new areas that support it and use it. It is a concise introduction to the most prominent, recent technological topics in computer networking. Topics include network technology such as wired and wireless networks, enabling technologies such as data centers, software defined networking, cloud and grid computing and applications such as networks on chips, space networking and network security. The accessible writing style and non-

mathematical treatment makes this a useful book for the student, network and communications engineer, computer scientist and IT professional.

Going beyond isolated research ideas and design experiences, Designing Network On-Chip Architectures in the Nanoscale Era covers the foundations and design methods of network on-chip (NoC) technology. The contributors draw on their own lessons learned to provide strong practical guidance on various design issues. Exploring the design process of the network, the first part of the book focuses on basic aspects of switch architecture and design, topology selection, and routing implementation. In the second part, contributors discuss their experiences in the industry, offering a roadmap to recent products. They describe Tilera’s TILE family of multicore processors, novel Intel products and research prototypes, and the TRIPS operand network (OPN). The last part reveals state-of-the-art solutions to hardware-related issues and explains how to efficiently implement the programming model at the network interface. In the appendix, the microarchitectural details of two switch architectures targeting multiprocessor system-on-chips (MPSoCs) and chip multiprocessors (CMPs) can be used as an experimental platform for running tests. A stepping stone to the evolution of future chip architectures, this volume provides a how-to guide for designers of current NoCs as well as designers involved with 2015 computing platforms. It cohesively brings together fundamental design issues, alternative design paradigms and techniques, and the main design tradeoffs—consistently focusing on topics most pertinent to real-world NoC designers.

New Scientist magazine was launched in 1956 "for all those men and women who are interested in scientific discovery, and in its industrial, commercial and social consequences". The brand's mission is no different today - for its consumers, New Scientist reports, explores and interprets the results of human endeavour set in the context of society and culture.