Tutorial One

Functional Verification of System on Chips - Practices, Issues and Challenges

Subir K. Roy (Coordinating Presenter)
Synplicity, Inc.,USA
subir@synplicity.com

S. Ramesh
Indian Institute of Technology Bombay, India
ramesh@cse.iitb.ernet.in

Supratik Chakraborty
Indian Institute of Technology Bombay, India
supratik@cse.iitb.ernet.in

Tsuneo Nakata
Fujitsu Laboratories Limited, Japan
nakata@flab.fujitsu.co.jp

Sreeranga P. Rajan
Fujitsu Laboratories of America, USA
sree@fla.fujitsu.com

Abstract

System on Chip (SoC) designs inherit all the well known verification and validation difficulties associated with complex ASIC designs, besides adding their own set of newer problems. These arise because SoCs are primarily implemented by re-using Intellectual Property (IP) cores. It is well known that verification today constitutes about 70% to 80% of the total design effort, thereby, making it the most expensive component in terms of cost and time, in the entire design flow. It is expected to get even worse for SoC designs.

In a complex SoC design flow functional verification is very important; any behavioral or functional bug escaping this phase will not be detected in the subsequent implementation phases and will surface only after the first silicon is integrated into the target system, resulting in costly design and silicon iterations. A number of academic and industrial research laboratories have been carrying out research on functional verification of SoCs based on different approaches. Partial success has been achieved in deploying them. Many of the issues relate to intrinsic limitations of some of the approaches taken; while others have to do with the quality of the design information, by way of, design descriptions, design documentations and design specifications, from which the overall verification objectives are derived. SoCs have brought to focus the need to carry out design and verification concurrently. For the design and verification task to proceed concurrently there is a need to capture formally, design information and implementation details at various levels of abstraction. Another reason for the need to formalize is that, as designs become more complex, functional verification will have to be carried out using the divide and conquer approach.

We discuss several approaches based on compositional verification. For these approaches to succeed, specifications of either, the individual modules, or individual IPs, if any are used, have to be stated formally. There exist several commercial offerings addressing the area of SoC functional verification.

Most of these use some form of divide and conquer approach related to compositional verification. The basis of success of some of these tools lies in the fact that the specifications of the IP cores are in essence captured in some executable form, be they formal specification languages, or commonly used HDLs suitably modified for the purpose.

This tutorial is structured to provide information on the state of the art in the area of functional verification. It will focus on existing methodologies, tools, and practical approaches based on universal simulation, emulation, formal verification, and semi-formal verification that can be employed to overcome the SoC verification problem. We discuss a number of real life verification projects, describing the various techniques used and the effectiveness of these techniques. We conclude the tutorial by presenting issues, which form the current focus for research.

Dr. Subir K. Roy *is currently a Project Leader at Synplicity Inc. Prior to that he was an Assistant Professor in the Department of Electrical Engineering, IIT Kanpur from Nov. 1993 to Jan. 2001. From Feb. 1998 to Jan. 2000 he worked as a Researcher in the CAD Laboratory of Fujitsu Laboratories Limited, Kawasaki, Japan, in the area of formal and semi-formal verification of VLSI chips. His contributions in this area are in the automatic formulation of certain classes of properties from the RTL description of designs, compositional verification of SoCs, and design abstraction for model checking.*

His broad areas of research interests are in the semi-formal & formal verification of SoC designs, Synthesis of Asynchronous Digital Systems, Low Power Synthesis and High Level Digital System Design. Dr. Subir K. Roy received the B.E. degree in Electronics & Telecommunication Engineering, from the University of Pune in 1982, the M.Tech. degree in Electrical Engineering from IIT Madras in 1984 and the Ph.D. degree in Electrical Engineering from IIT Bombay, in 1993. Dr. Roy has a number of publications, and has delivered several invited talks and tutorials in academic and commercial organizations.

Dr. S. Ramesh *has more than 10 years of research experience in the areas of formal specification, verification and design of concurrent and reactive languages. He received the BE degree in electronic and communicaion engineering from the Indian Institute of Science, Bangalore, India in 1981, and the PhD degree in Computer Science and Engineering from Indian Institute of Technology, Bombay, India in 1987. From Aug. 1987 till Dec. 1988, he was a post-doctoral fellow at Technical University Eindhoven, Netherlands. He was then a visiting researcher at Tata Institute of Fundamental Research, Bombay for 9 months. His collaborative research with Gerard Berry's group at INRIA, Sophia-Antipolis, France, led to the development of a language for describing distributed controllers called Communicating Reactive Processes (CRP). Currently he is a Full Professor in the Dept of Computer Science and Engineering, Indian Institute of Technology, Bombay. He is also heading the Centre for Formal Design and Verification of Software recently set up in IIT Bombay to develop tools and techniques for formal verification of industrial software and hardware systems.*

Dr. Supratik Chakraborty *received the B.Tech. in Computer Science and Engineering from Indian Institute of Technology (IIT), Kharagpur, and subsequently received the MS and Ph.D. degrees in Electrical Engineering from Stanford University. He has served for one year as Member of Research Staff in the Advanced CAD Research group at Fujitsu Labs of America, Inc.*

Currently, he is an Assistant Professor in Computer Science and Engineering at IIT Bombay, and is a Principal Investigator in the Centre for Formal Design and Verification of Software at IIT Bombay. His primary research interests are in the use of formal methods for verification and analysis of digital

systems. He is also interested in complexity and approximability issues in verification problems. He has also worked on approximation techniques for timing analysis and verification of asynchronous systems.

He has taught a post-graduate level course on asynchronous system design and verification at IIT Bombay. He has also given a tutorial lecture on BDDs and hardware verification at a workshop on formal methods held at IIT Bombay. In addition, he has given several invited presentations on formal verification methods, asynchronous systems and timing analysis at academic and commercial organizations.

Dr. T. Nakata *received the B.S. degree in Electronic Engineering and the M.S. and Ph.D. degrees in Information Engineering from the University of Tokyo, Tokyo, Japan in 1981, 1983 and 1986, respectively. He joined Fujitsu Laboratories Ltd (FLL), Kawasaki, Japan in 1986 and has been engaged in research and development of computer aided design. He spent one year as a Visiting Scholar at the University of California at Berkeley from 1993 to 1994. He is a member of the IEEE and the Information Processing Society of Japan (IPSJ). His research interests include design methodology and verification of VLSIs. He presently heads the System Level Verification group at FLL.*

Dr. Sreeranga P. Rajan *is currently a Member of Technical Staff at Fujitsu Research Laboratories (Sunnyvale, California). Dr. Rajan's research contributions span from, developing methods for formal verification and system design, to developing compilers for embedded processors. Dr. Rajan's research in high-level verification led to the development of an automatic high-level model checking tool that has been used to debug large designs such as an ATM switch and to debug high-level synthesis tools. Dr. Rajan's early work led to the first system that integrates two powerful techniques in formal verification: theorem-proving and model-checking. The integration, which consists of efficient computation schemes within a theorem-proving framework, has led to automatic strategies for solving hard/large verification problems in software and hardware. The integrated system is distributed world-wide and has been used in industries and universities in Projects, including network protocol design, software requirements analysis, and hardware verification. Dr. Rajan has numerous publications, and has given several invited talks and tutorials.*

Tutorial Two

System-Level Design of Embedded Media Systems

Pieter van der Wolf (Coordinating Presenter)
Philips Research Laboratories, The Netherlands
Pieter.van.der.Wolf@philips.com

Wido Kruijtzer
Philips Research Laboratories, The Netherlands
Wido.Kruijtzer@philips.com

Jos van Eijndhoven
Philips Research Laboratories, The Netherlands
Jos.van.Eijndhoven@philips.com

Abstract

This tutorial reports on system-level design practices in Philips. First we present the design problems encountered in the development of embedded media systems for the consumer market. The characteristics and requirements of the consumer electronics domain are presented. We focus on high performance video applications and the demands that these applications put on architectures of embedded media systems.

Several high performance embedded systems architectures are illustrated. This includes an in depth discussion of high performance programmable components, in which the TriMedia VLIW cores are taken as an example. We further discuss system-on-a-chip (SoC) architectures that incorporate such cores. We address the aspect of heterogeneous SoC architectures that incorporate components in the range from programmable processors to function-specific hardware engines. For such architectures trade-offs can be made with respect to flexibility, performance, and cost to implement an optimal solution. This process involves hardware/software co-design with evaluation of different points in the architecture design space. A key aspect of SoCs is the communication architecture by means of which the different components communicate. We discuss the key issues in the design of such communication architectures. SoC architectures need to include a software stack and related programming paradigm to allow them to be programmed efficiently. We present the key issues that need to be addressed by an SoC software stack and illustrate this with an example.

The design of advanced SoCs and related software is a formidable task. A recent trend in this respect is platform-based design. Platform-based design aims to support the development of advanced SoCs by providing an architecture template, or platform, for the design of such SoCs. Platform-based design aims to facilitate reuse of hardware and software components and supports the rapid development of derivative designs. It thereby enables cost-effective design of a family of products under stringent time-to-market constraints. We present the essential ingredients of platform-based design and illustrate these with real-life examples.

Advanced design technology is required to tackle the design problems in the design of embedded media systems. We present a number of key principles for quantitative and explorative design and discuss several levels of abstraction at which designs can be modeled and evaluated. We further elaborate on the development and application of design technology based on a number of design cases. These cases

include the development of a TriMedia 64-bit VLIW CPU and heterogeneous multi-processor architectures for media processing.

The tutorial includes relevant references to related work in the field of embedded media systems and related design technology. The tutorial concludes with a summary that identifies a number of key lessons learned.

Pieter van der Wolf *was born May 20, 1961, in Gouda, The Netherlands. He obtained both an MSc degree (1986) and a PhD degree (1993) in Electrical Engineering from Delft University of Technology, The Netherlands. His PhD research concerned the definition and implementation of CAD Framework technology. After obtaining his PhD degree he became an assistant professor at Delft University of Technology. In 1996 he joined Philips Electronics, to become Senior Scientist at Philips Research Laboratories in Eindhoven, The Netherlands. He is active in the field of embedded systems architectures, with a focus on design technology. His research interests include system-level design methodologies, embedded systems architectures, and embedded processors.*

Wido Kruijtzer *was born October 22, 1965, in Roermond, The Netherlands. He holds a BSc degree in Electrical Engineering from Heerlen University, The Netherlands, and a MSc degree in Electrical Engineering from Eindhoven University of Technology, The Netherlands. He currently is a Senior Scientist at Philips Research Laboratories in Eindhoven, The Netherlands. His research field is in the area of embedded digital systems and includes design technology and architectures. He regularly serves on program committees of major conferences in his field, such as DATE.*

Jos T. J. van Eijndhoven *was born March 1, 1957, in Roosendaal, The Netherlands. He studied electrical engineering at the Eindhoven University of Technology, The Netherlands, obtaining his MSc degree in 1981 and his PhD degree in 1984, for his work on piecewise linear circuit simulation. He became senior research member in the design automation group of the Eindhoven University of Technology, responsible for covering the system and architectural synthesis research areas. In 1986 he spent a sabbatical period at the IBM Thomas J. Watson Research Laboratory, Yorktown Heights, New York, for research on high level synthesis. In 1998 he moved to Philips Research Laboratories in Eindhoven, to work on the architectural design of programmable multimedia hardware and the associated mapping of media processing applications.*

Tutorial Three

Trends and Challenges in VLSI Technology Scaling Towards 100nm

Stefan Rusu (Coordinating Presenter)
Intel Corporation, USA
stefan.rusu@intel.com

Manoj Sachdev
University of Waterloo, Canada
msachdev@ece.uwaterloo.ca

Christer Svensson
Linkoping University, Sweden
chs@isy.liu.se

Bram Nauta
University of Twente, The Netherlands
B.Nauta@el.utwente.nl

Abstract

Moore's Law drives VLSI technology to continuous increases in transistor densities and higher clock frequencies. This tutorial will review the trends in VLSI technology scaling in the last few years and discuss the challenges facing process and circuit engineers in the 100nm generation and beyond. The first focus area is the process technology, including transistor scaling trends and research activities for the 100nm technology node and beyond. The transistor leakage and interconnect RC delays will continue to increase.

The tutorial will review new circuit design techniques for emerging process technologies, including dual Vt transistors and silicon-on-insulator. It will also cover circuit and layout techniques to reduce clock distribution skew and jitter, model and reduce transistor leakage and improve the electrical performance of flip-chip packages.

Another focus area is the circuit design and the techniques used to minimize capacitive and inductive noise. The tutorial will review models of electrical interconnects, including inductance and skin effect. These models are used to estimate the performance of electrical interconnects, including delays, data-rates and power consumption for on-chip and off-chip interconnects and for clock distribution.

Integrating analog circuits on large digital chips presents significant challenges, primarily due to substrate noise coupling. This tutorial will describe a strategy for making analog circuits less sensitive to substrate bounce, including examples for epi-type CMOS technology.

Finally, the tutorial will review the test challenges for the 100nm technology node due to increased clock frequency and power consumption (both active and passive) and present several potential solutions.

Stefan Rusu *is a Principal Engineer in Intel's Enterprise Products Group leading the technology and special circuits design group for all the Itanium Processor Family designs. He first joined Intel Corp. in 1984 working on data communications integrated circuits. In 1988 he joined Sun Microsystems working*

on microprocessor design with focus on clock and power distribution, packaging, standard cell libraries, CAD and circuit design methodologies. He re-joined Intel Corp. in 1996 working on the clock and power distribution, cell library, I/O buffers and packaging of the first Itanium microprocessor. He is presently developing circuit design methodologies for Intel's 100nm process generation. He received the MSEE degree from the Polytechnic Institute in Bucharest, Romania. He has published numerous technical papers and has been an invited speaker at several conferences. Stefan currently holds 12 U.S. patents with several more pending. He is a Senior Member of IEEE and has been a member of the ESSCIRC Technical Program Committee since 1998.

***Manoj Sachdev** is an associate professor in the electrical and computer engineering department at University of Waterloo, Canada. His research interests include low power and high performance digital circuit design, test and manufacturing issues of integrated circuits. He has written a book, two book chapters on testing and has published significantly in conferences and journals. He received best paper award for his paper in European Design and Test Conference, 1997 and an honorable mention award for his paper in International Test Conference, 1998.*

***Christer Svensso**n is professor in Electronic Devices at Linkoping University. He was born in Boras, Sweden in 1941 and received the M.S. and Ph.D. degrees from Chalmers University of Technology, Sweden, in 1965 and 1970 respectively. He was with Chalmers University from 1965 to 1978, where he performed research on MOS transistors, nonvolatile memories and gas sensors. He joined Linkoping University 1978, and is since 1983 professor in Electronic Devices there. He initiated a new research group on integrated circuit design. Svensson's present interests are high performance and low power analog and digital CMOS circuit techniques for communication, computing and sensors. Svensson has published more than 160 papers in international journals and conferences and holds 8 patents. He was awarded the Solid-State Circuits Council 1988-89 best paper award. He is a member of the Royal Swedish Academy of Engineering Sciences. He is a cofounder of several companies, most recently Switchcore AB, Optillion AB and Bluetronics AB.*

***Bram Nauta** was born in Hengelo, The Netherlands, in 1964. In 1987 he received the M.Sc. degree (cum laude) in Electrical Engineering from the University of Twente, Enschede, The Netherlands. In 1991 he received the Ph.D. degree from the same university on the subject of analog CMOS filters for very high frequencies. In 1991 he joined the Mixed-Signal Circuits and Systems Department of Philips Research, Eindhoven, The Netherlands, where he worked on high speed AD converters. >From 1994 he led a research group in the same department, working on "analog key modules". In 1998 he returned to the University of Twente, as full professor heading the IC Design group in the MESA+ Research Institute and department of Electrical Engineering. His current research interest is analog CMOS circuits for transceivers. He is also part-time consultant in industry and in 2001 he co-founded Chip Design Works. His Ph.D. thesis was published as a book: Analog CMOS Filters for Very High Frequencies, Kluwer, Boston, MA, 1993. He holds 8 patents in circuit design and he received the "Shell Study Tour Award" for his Ph.D. Work. From 1997-1999 he served as Associate Editor of IEEE Transactions on Circuits and Systems -II; Analog and Digital Signal Processing, and in 1998 he served as Guest Editor for IEEE Journal of Solid-State Circuits. In 2001 he became Associate Editor for IEEE Journal of Solid-State Circuits.*

Tutorial Four

Mathematical Methods in VLSI

M. V. Atre (Coordinating Presenter)
Agere Systems (Lucent Technologies), Bangalore, India
mvatre@agere.com

P. S. Subramanian
Sasken Communication Technologies Limited, Bangalore, India
pss@sasken.com

H. Narayanan
Indian Institute of Technology, Mumbai, India
hn@ee.iitb.ernet.in

Abstract

The theme of the tutorial is the use of mathematical methods in VLSI. The traditional use of mathematics in engineering disciplines is via mathematical modeling- concepts and interactions in the problem domain are mapped to objects and relationships of a specific mathematical topic and then the formal deductions within the topic are re-interpreted in the problem domain. After a structured review of VLSI design flow and the identification of mathematical topics applicable to each step of the design flow, the tutorial illustrates these themes by a sampling of mathematical techniques applicable to analysis of modeling and simulation, partitioning, structural and behavioral decomposition, and symbolic reasoning about behavior.

The tutorial is aimed at illustrating the importance of mathematics in VLSI, especially in the development of various tools, which are critical for design. The audience will be made to appreciate how some of the tools which are used by designers actually have some very deep mathematics built into them, without which it would be impossible for any automation in the design process. This mathematics becomes more important as we go to high complexity designs involving millions of transistors, high frequencies and systems-on-chip.

The morning session will have a brief introduction on what will be covered in the tutorial, and why. This will be followed by a brief description of the existing design flow concepts, as well as how the various areas of mathematics interleave with the various aspects of design flow. There will be brief presentations and overviews of some mathematical topics which are relevant from the VLSI perspective: a) the uses of graphs and hypergraphs for modeling real life problems particularly in terms of connection, flow of information, and hierarchical construction; b) some samples of the use linear algebra in VLSI and in particular the use of eigenvalues in various contexts such as system reduction and partitioning; c) introduction to submodular functions; d) the importance of differential equations; e) classification of differential equations and some of the important ways to solve them; f) the foundations of symbolic methods used in modeling behavior for analysis and verification g) the ubiquitous `decomposition theme'.

The afternoon session will take up a few areas of mathematics in greater detail, how they help address some the very important areas in VLSI design automation, and how they can possibly address future challenges: a) use of submodular functions for optimization problems in VLSI, their use for partitioning large scale systems which arise in many forms throughout this area is described (including the realization of large scale finite state machines through decomposition techniques and the parallel simulation of large electrical circuits); b) the use of differential equations in modeling and simulation, the SPICE equations,

their solutions using various techniques, and the relation between SPICE simulation and various other simulation scenarios; c) logical and algebraic techniques in verification d) applications of classical automaton decomposition theory.

At the end of the tutorial, we expect the audience to have a better appreciation of how mathematics is important, relevant and critical for developing design automation tools.

Dr. Madhusudan V. Atre *did his 5 year integrated MSc in Physics from IIT-Bombay, with specialisation in Solid State Physics in 1978. He obtained his PhD in Theoretical Physics from IISc-Bangalore in 1985, specialising in Nonlinear Phenomena in Plasma Physics, Low Temperature Physics, and Mathematical Methods in Dynamical Systems. From 1986-1990, he was Research Associate at the TIFR-Bombay, PRL-Ahmedabad, and IISc-Bangalore, working in the areas of Mathematics of Quantum Field Theory and Dynamical Systems. In 1990 he joined the DRDO where he headed the Software Group, looking at scientific applications on parallel computers, VLSI test generation techniques, Lie-group symmetries of differential equations. He was manager of VLSI activities in Crosscheck Tech, before joining Texas Instruments in 1995 where he headed the Design Flow and Systems Engineering team, as well as the core-EDA activities. In 1998, he started the Bell Labs R&D center of the Lucent Technologies, Microelectronics Group, (now called Agere Systems after spinoff from Lucent Technologies) . As the Managing Director, he coordinates research and development in the areas of VLSI and Optoelectronics CAD, DSP software tools and applications, as well as VLSI design.*

Mr. P. S. Subramanian *obtained his degree in Electrical Engineering from VJTI, Mumbai in 1968. He joined TIFR, Mumbai in 1969 and worked there till 1999. He then joined Sasken Communication Technologies and is currently heading their Computer Science R&D Group working on Adaptive Embedded Systems. During his career he has worked in the areas of VLSI CAD and Computer Science ranging from the very practical to the highly theoretical. Initially, he worked in Graphics and Time Sharing Software. He then built a scientific calculator using MSIs by first designing an 8-bit processor and then microcoding CORDIC algorithms into it. He developed and made extensive use of CAD tools during this development. With this background he obtained an UNDP fellowship to study CAD in the LOCMOS facility of Philips at Nijmegen, Holland. After his return he directed a team developing a suite of CAD tools with grants from DoE. He then played an active role in a task force set up by DoE to promote VLSI activities in India. Later, he worked in areas of theoretical computer science ranging from complexity theory to semantics. He is interested in reusing the "methodologies" developed in the domain of mathematics for system design.*

Prof. H. Narayanan *did his B.Tech in Electrical Engineering in 1969 and his PhD in Electrical Engineering in 1974, both from I.I.T. Bombay. Since 1974 he has been with the EE Dept., I.I.T. Bombay as a faculty member and is currently a Professor. During 1983-85 he was Visiting Associate Professor at EECS, U.C. Berkeley. H. Narayanan's interest is in the area of network topology in the context of efficient network analysis. Towards this end he has worked at understanding the interactions of graph theory, matroid theory and submodular function theory with Electrical network theory. His main results in Electrical networks include a topological theory of network decomposition. In the area of matroids and submodular functions his results include the first complete description of the principal partition of a matroid (1974)(independent of N. Tomizawa of Tokyo University [1976]) and the theory of the principal lattice of partitions of a submodular function (1986). He is the author of the monograph "Submodular Functions and Electrical Networks" (Annals of Discrete Maths vol. 54). He has supervised the building of the general purpose circuit simulator BITSIM (based on topological hybrid analysis) and the point relaxation based circuit simulator BREMICS at the VLSI Design Centre, I.I.T. Bombay.*

Tutorial Five

Electronic Testing for SOC Designers

Vishwani D. Agrawal (Coordinating Presenter)
Agere Systems, USA
va@agere.com

Michael L. Bushnell
Rutgers University, USA
bushnell@caip.rutgers.edu

Abstract

The speakers have given VLSI design and VLSI testing courses at Rutgers University for over ten years. They combine their industry and academic experiences, and believe that today's VLSI designer must know the essentials of digital, memory and mixed-signal circuits, as well as the test standards for systems that combine such components. This tutorial provides a careful selection of topics on testing of all three types of circuits and systems. The presentation is divided into three parts. Part I (Introduction) contains definition of test and its motivation, test process and automatic test equipment (ATE), test economics and product quality, and fault modeling. Part II (Test Methods) includes logic and fault simulation, testability measures, combinational and sequential ATPG, memory test, DSP-based analog test, model-based analog test, delay test, and IDDQ test. Part III (Design for Testability) covers scan design, built-in self-test (BIST), boundary scan, analog test bus, system test and testing of core-based designs. The tutorial is derived from the authors' recent textbook, *Essentials of Electronic Testing for Digital, Memory and Mixed-Signal VLSI Circuits* (Kluwer Academic Publishers, Boston, 2000, ISBN 0-7923-7991-8) and their full-semester course material available at the website http://cm.bell-labs.com/cm/cs/who/va.

Vishwani D. Agrawal *is a Distinguished Member of Technical Staff in the Circuits and Systems Research Lab of Agere Systems (a microelectronics company to be spun off from Lucent Tech.) He holds a PhD from the University of Illinois. He is a Fellow of the IEEE. He has worked in the area of electronic testing for 30 years, has published 250 papers and 5 books and holds 13 U.S. patents.*

Michael L. Bushnell *is a Professor in the Department of Electrical and Computer Engineering at Rutgers University. He holds a PhD from Carnegie Mellon University. He has been a researcher in electronic testing and other VLSI related areas for over 15 years. He has published 4 books and numerous papers. He currently serves as the Graduate Director of his department at Rutgers University.*

Tutorial Six

Specification, Modeling and Design Tools for System-on-Chip

Luciano Lavagno (Coordinating Presenter)
Cadence Berkeley Laboratories, USA
luciano@cadence.com

Sujit Dey
University of California, San Diego, USA
dey@ece.ucsd.edu

Rajesh Gupta
University of California, Irvine, USA
gupta@uci.edu

Abstract

Ubiquitous embedded systems are revolutionizing our daily lives. Whole systems on a chip deliver unprecedented computation power at ever decreasing costs. However, their complexity makes their design with traditional RTL-based flows extremely challenging. Complexity in such systems arises not only from the diversity of the technologies, from RF front-ends to baseband DSP software, that must be integrated on-chip, but also from the fact that such systems must be increasingly built from parts that have been designed separately and using different tools and flows. High abstraction levels and component re-use are essential. Often such systems are highly networked and rely on sophisticated communication mechanisms. Architectural design and performance analysis of such networked system-chips is a crucial part of the embedded system design process.

Two basic methods for tackling the growing complexity of system-on-chip design are emerging. Formal specification and design methods, and platform-based design. Formal methods are essential to capture precisely the designer intent at the highest possible level of abstraction, and to guide him down towards an implementation. Design errors are reduced by the use of synthesis techniques, as well as by the availability of a variety of validation techniques. Moreover, the orthogonalization of concerns (e.g., computation vs. communication, function vs. architecture) promotes extensive reuse of both parts of an application (derivative designs) and of previously used architectures.

A platform in this context is the precise specification of a high-level implementation target that supports both a variety of lower level implementations, and a variety of higher-level applications. A classical example is the PC, that has supported an explosive growth of both applications and implementations. In embedded systems-on-chip platform clearly need to be application-specific. A typical platform includes one or more processors, memory, one or more bus hierarchy levels, and peripherals devoted to application-specific functions (e.g., MPEG decoding, filtering, A/D conversion, ...). Mapping to a platform involves identifying which portions of the application will be implemented on each computation and communication resource. Communication must be refined from high-level primitives (e.g., FIFOs) down to platform primitives (e.g., interrupts, memory buffers, ...). Functionality must be synthesized and scheduled based on platform-specific cost, memory and performance constraints.

In this tutorial we will cover both aspects in detail, illustrating first the languages and Models of Computation available to the system-level designer to capture precisely and unambiguously the

requirements. We will discuss for what application domain and platform each language is most appropriate, focusing mostly on platform-independent languages and MOCs, since they support the greatest freedom in mapping choice.

We will then discuss how the architecture of the platform, and the services it offers to the application designer, can also be formally and compactly captured and specified. We will show how the mapping paradigm can be used to select an implementation for the functional blocks and their communication, and how simulation and implementation methods can be derived automatically. In particular, we will describe how software estimation and synthesis for reactive real-time systems can be competitive with hand design, while retaining the ease of re-use typical of high-level specifications.

In the next part of the tutorial, we will establish the importance of on-chip communication architectures in determining the performance of System-on-Chips (SoCs). We will track several on-chip communication architectures that are in use today, including different types of bus architectures, token ring architectures, and crossbar switches. We will analyze the performance of the existing communication architectures under different classes of on-chip communication traffic. We will also describe novel on-chip communication architectures that are being developed to satisfy the growing needs of fast and concurrent on-chip communication in very high-performance systems like network processors.

Next, we will describe efficient methodologies that can be used to analyze and design customized on-chip communication architectures for new application-specific platforms, as well as optimally map an application's communication requirements to the communication architecture present in an existing platform. We will also discuss reconfigurable on-chip communication architectures that allow runtime adaptation of the communication protocols to changing communication demands of the platform. Besides performance, we will analyze the energy consumed by on-chip communication architectures, and describe techniques to minimize such energy consumption. Lastly, we will analyze the impact of deep sub-micron technologies on SoC communication, and discuss the needs for development of noise-aware on-chip communication.

Luciano Lavagno *graduated in Electrical Engineering from Politecnico di Torino (Italy) in 1983. From 1984 to 1988 he was with CSELT Laboratories (Torino, Italy). In 1992 he received his Ph.D. in Electrical Engineering and Computer Science from the University of California at Berkeley. Between 1993 and 1998 he was an Assistant Professor with the Department of Electronics of Politecnico di Torino. Since 1993 he has been the architect of the POLIS project developing a complete hardare/software co-design environment for control-dominated embedded systems. Since 1998 he has been an Associate Professor with the Department of Electrical, Management and Mechanical Engineering (DIEGM) of the University of Udine, Italy. His research interests include the synthesis of asynchronous circuits, and the concurrent design of mixed hardware and software systems.*

Dr. Lavagno is the author of a book on asynchronous circuit design, the co-author of a book on hardware/software co-design of embedded systems, and has published over 100 journal and conference papers. He has served as technical committee member of several international conferences in his field (namely the Design Automation Conference, the International Conference on Computer Aided Design, the conference on Design Automation and Test in Europe) and as technical committee member or chair of several workshops and symposia.

Sujit Dey *received the Ph.D. degree in Computer Science from Duke University in 1991. From 1991 to 1997, he was at the NEC C&C Research Laboratories, Princeton, NJ, where he was a Senior Research Staff Member. While at NEC, he developed several design methods and tools for the design of high performance and low power system-on-chips, leading to several technology transfers both within NEC as well as outside. In 1998, he joined the University of California, San Diego, where he is an Associate*

Professor in the Electrical and Computer Engineering Department. His research group at UCSD is developing configurable platform architectures, consisting of adaptive wireless applications, protocols, and hardware-software architectures for next-generation wireless network appliances. He is affiliated with the California Institute of Telecommunications and Information Technology, the UCSD Center for Wireless Communications, and the DARPA/MARCO Gigascale Silicon Research Center. Dr. Dey has co-authored more than 100 publications, including journal and conference papers, a book and several book chapters. He received Best Paper awards at the Design Automation Conferences in 1994, 1999, and 2000, and the 11th VLSI Design Conference in 1998, and several best paper nominations. He is a co-inventor of 8 U.S. patents, and has several others pending. He has presented numerous full-day and embedded tutorials, and participated in panels, in the topics of low-power wireless systems design, hardware-software embedded systems, and deep sub-micron system-on-chip design and test. He has been the General Chair and Program Chair, and member of organizing and program committees, of several IEEE conferences and workshops.

Rajesh Gupta *(Ph. D. Stanford, M.S. UC Berkeley) is an associate professor in Information and Computer Science at University of California, Irvine. His research interests are in system-level design for embedded and portable systems, VLSI design, and adaptive system architectures. He worked as an assistant professor at University of Illinois, Urbana-Champaign from 1994 through 1996. Prior to that he was at Intel Corporation in Santa Clara, California where he worked as a member on a number of processor design teams. He is co-author of three patents on PLL-based clock circuit; synthesis with regularity and system-on-chip modeling and a patent (filed) on data-path synthesis. He is author of a book on* Co-synthesis of Hardware and Software for Digital Embedded Systems *published by Kluwer Academic in 1995. At UCI, he leads an effort on Adaptive Memory System Architectures and co-leads an effort on Compiler-Controlled Power/Performance Management both sponsored by DARPA programs. Gupta serves as Chair of the CANDE technical committee and as a board of governor of the IEEE Circuits and Systems Society. He also serves as associate editor-in-chief of IEEE Design and Test and on the editorial board of IEEE Transactions on CAD.*

Tutorial Seven

MEMS: Technology, Design, CAD and Applications

R. Lal (Coordinating Presenter)
Indian Institute of Technology, Bombay, India
rlal@ee.iitb.ac.in

Prakash R Apte
Tata Institute of Fundamantal Research, Mumbai, India
apte@tifr.res.in

K. N. Bhat
Indian Institute of Technology, Madras, India
knbhat@ee.iitm.ernet.in

G. Bose
Indian Institute of Technology, Delhi, India
gbose@care.iitd.ernet.in

S. Chandra
Indian Institute of Technology, Delhi, India
schandra@care.iitd.ernet.in

D. K. Sharma
Indian Institute of Technology, Bombay, India
dinesh@ee.iitb.ac.in

Abstract

Micromachined Electro-Mechanical Systems(MEMS), also called Microfabricated Systems (MS), have evoked great interest in the scientific and engineering communities. This is primarily due to several substantive advantages that MEMS offer: orders of magnitude smaller size, better performance than other solutions, possibilities for batch fabrication and cost-effective integration with electronics, virtually zero dc power consumption and potentially large reduction in power consumption, etc. The application domains cover microsensors and actuators for physical quantities (MEMS), of which MEMS for automobile & consumer electronics forms a large segment; microfabricated subsystems for communications and computer systems (RF-MEMS & MOMS); and microfabricated systems for chemical assay (microTAS) and for biochemical and biomedical assay (bioMEMS and DNA chips). This tutorial would give an introduction to these exciting developments and the technology and design approaches for the realization of these integrated systems.

The full day tutorial would begin with a synoptic overview of the area, highlight some of the challenges and outline the scope of the tutorial. It would be followed with an introduction to the design of microsensors, such as the pressure sensor and the accelerometer, that began the MEMS revolution. We begin with a quick introduction to material properties at the micron scale and show that silicon is eminently suited for micromechanical devices and therefore the possibility of integrating MEMS with

VLSI electronics. Unit processes for bulk and surface micromachining of silicon and integration of processes for fabricating silicon microsensors will be presented.

Smart cell phones and wireless enabled devices are poised to become commercial engines for the next generation of MEMS, since MEMS provide not only better functionality with smaller chip area, but also alternative transceiver architectures for improved functionality, performance and reliability. We, therefore, have two lectures on the structure and design of RF-MEMS covering passive & resonant elements, switches and transmission lines. Some applications of MOMS in communication systems will also be discussed.

We shall have a lecture on bioMEMS to highlight the immense possibilities that exist for MEMS in the life sciences & medicine. The idea of integrating microfluidics and biological or biomimetic material with electronic systems is alien to electronic systems designers and there are problems with integrating wet systems with electronics. We give a synopsis of the types of structures required and approaches for the design and test of such systems.

Finally, we shall discuss the issues involved with embedding MEMS in complete systems, including issues related to design tools, simulation, test and parameter extraction & de-embedding.

The faculty for the tutorial have been in the forefront of integrated circuit technology development in the country and are currently involved with various facets of research with MEMS.

Dr. Lal *is a Professor of Electrical Engineering at IIT Bombay. His research interests include the physics and modelling of semiconductor devices, radiation and high-field effects in devices and circuits, and device characterization. He has also been working on radiation sensors, biosensors and biosensing systems, much of the latter work as part of an interdepartmental effort involving Chemistry, Materials Science and Biomedical Engineering.*

Dr. Apte *is at TIFR and would be a visiting professor at IIT Bombay, slated late autumn 2001. He had been instrumental in developing India's first TTL technology. Current, research interests are micromachining, high temperature superconductor devices and quality improvement using innovative problem solving tools.*

Dr. Bhat *is a Professor of Electrical Engineering at IIT Madras and has coordinated the microelectronics activity of the department. His research interests include SOI MOSFET modeling and technology, GaAs and InP surface passivation, GaAs MISFET technology and modeling, polysilicon thin film transistors and grain-boundary passivation, and silicon icromachining and micromachined sensors.*

Dr. Bose *and* ***Dr. Chandra*** *are faculty members at CARE, IIT Delhi and have both contributed substantially to CARE's microelectronics development. Dr. Bose has worked extensively with lithography & CMOS process integration and is currently interested in micromachined sensors for temperature, pressure and sound. Dr. Chandra has worked in the areas of plasma CVD and RTP processes, laser recrystallization, direct wafer bonding, and SIMOX based SOI technologies. Current interests are micromachined sensors and actuators with focus on switches and millimeter wave devices using MEMS.*

Dr. Sharma *is a Professor of Electrical Engineering at IIT Bombay. His interests are broadly in the areas of MOS device modeling, simulation & characterization, and VLSI design & technology. His current interests include mixed signal and digital signal processing VLSIs, asynchronous design, embedded system design, radiation hard technology and the effect of technology and device scaling on design architectures and tools.*

Tutorial Eight

Logic Design of Asynchronous Circuits

Jordi Cortadella (Coordinating Presenter)
Universitat Politecnica de Catalunya, Spain
jordic@lsi.upc.es

Alex Yakovlev
University of Newcastle upon Tyne, United Kingdom
alex.yakovlev@ncl.ac.uk

Jim Garside
University of Manchester, United Kingdom
jgarside@cs.man.ac.uk

Abstract

This tutorial aims at motivating the audience to consider asynchronous circuits as a competitive alternative to solve some of the design problems inherent to submicron technologies. One of the main reasons why designers are reluctant to incorporate asynchrony in their systems is the difficulty to design asynchronous circuits.

Asynchronous circuits are promising to tackle problems such as electro-magnetic interference, power consumption, performance, and modularity of digital circuits. The tutorial will introduce state-of-the-art tools and methodologies for their design. It will cover aspects such as specification, architectural design and controller synthesis tools, of asynchronous circuits. The tutorial will concentrate on a particular design methodology for control circuits based on specifications with Signal Transition Graphs. It will also cover design strategies for the microarchitecture, data-path and control circuits that have been successfully applied in the design of the asynchronous version of the ARM microprocessor.

The tutorial is organized in three parts. The first part will give an overview of the main concepts related to asynchrony in circuit design: asynchronous communication, design styles, building blocks, control specification and implementation, delay models, classes of asynchronous circuits and advantages of asynchronous circuits.

The second part will focus on the synthesis of control circuits. After a brief presentation of different specification formalisms, this part will describe a design flow based on specifications given as timing diagrams (Signal Transition Graphs). This design flow can be fully automated and the steps involved in this automation will be described: state encoding, derivation of next-state equations for speed-independent circuits, hazard-free logic decomposition and synthesis with relative timing. Some basic concepts on testing and formal verification will also be covered.

The third part is devoted to present real experiences in the design of asynchronous circuits. In particular, the asynchronous ARM microprocessor is used to illustrate the design of relevant parts of an asynchronous system (data-path, control, memory). This part will cover both architectural and design issues that are relevant to asynchronous design: conditional execution, multi-cycle instructions, "elastic" pipelines, self-timed memories, asynchronous caches, asynchronous intra-chip communication and globally-asynchronous locally-synchronous systems.

The presenters of the tutorial have a wide experience in the design methodologies of asynchronous circuits. All of them are members of the Working Group on Asynchronous Circuit Design (ACiD-WG), a consortium funded by the European Commission to promote RTD activities around the theme of asynchronous circuit design (see http://www.scism.sbu.ac.uk/ccsv/ACiD-WG for more details).

Jordi Cortadella *received the M.S. and Ph.D. degrees in Computer Science from the Universitat Politecnica de Catalunya, Barcelona, Spain, in 1985 and 1987 respectively. He is a Professor at the Department of Software of the Universitat Politecnica de Catalunya. In 1988, he was a Visiting Scholar at the University of California, Berkeley. His research interests include computer-aided design of VLSI systems with special emphasis on synthesis and verification of asynchronous circuits, concurrent systems and HW/SW co-design. He has coauthored over 100 research papers in technical journals and conferences. He has served on the technical committees of several international conferences in the field of Design Automation and Concurrent Systems. He was the Symposium Co-Chair of the 5th International Symposium on Advanced Research in Asynchronous Circuits and Systems in Barcelona, 1999. He is the main author of the tool petrify (www.lsi.upc.es/~jordic/petrify), currently used by several industries and Universities for the synthesis of asynchronous control circuits.*

Alex Yakovlev *received the MSc (1979) and PhD (1982) degrees in Computing Science from Electrotechnical University of St. Petersburg, Russia, where he worked in the area of asynchronous and concurrent systems since 1980, and in the period between 1982 and 1990 held positions of Assistant and Associate Professor at the Computing Science department. Since 1991 he has been a Lecturer, Reader and from 2000 Professor in Computer Systems Design at the Newcastle University Department of Computing Science, where he is heading the VLSI Design research group. His current interests and publications are in the field of modelling and design of asynchronous, concurrent, real-time and dependable systems. He has coauthored over 100 research papers in technical journals and conferences. He has served on the technical committees of several international conferences in the field of Asynchronous Systems, Concurrency and Petri nets. He was the Programme Committee Co-Chair of the 5th International Symposium on Advanced Research in Asynchronous Circuits and Systems in Barcelona, 1999, the Co-Organiser of the IEEE Workshop on Asynchronous Interfaces, Delft, 2000, and the Co-Organiser of two workshops and an advanced tutorial on "Hardware Design and Petri Nets", Lisbon (1998), Williamsburg (1999) and Aarhus (2000).*

Jim Garside *gained a BSc in Physics in 1983 at the University of Manchester and a PhD in Computer Science at the same institution in 1987. Since that time he has worked in hardware systems using Inmos Transputers for investigation into parallel computer architectures and as a programmer on Air Traffic Control systems. He was appointed as a lecturer in the Department of Computer Science at the University of Manchester in 1991. Since that time he has primarily been involved in research into asynchronous logic systems, especially in the design of the AMULET series of asynchronous microprocessors. Most recently he led the design effort on the AMULET3i asynchronous system-on-chip which includes the first asynchronous microprocessor to achieve parity in performance and power-efficiency with its synchronous direct equivalent.*