A From MoC to SoC – Programming Embedded Multiprocessor Systems

Organisers: Rainer Leupers, RWTH Aachen U, DE
Gerhard Fettweis, TU Dresden, DE

Speakers: Alan Gatherer, Huawei Technologies, US
Gerhard Fettweis, TU Dresden, DE
Marco Bekooij, NXP, NL
Jos van Eijndhoven, Vector Fabrics, NL
Rainer Leupers, RWTH Aachen U, DE
Matthias Weiss, Intel, DE

The trend towards multicore, and even manycore, hardware platforms is a great revolution in the embedded computing industry. In particular, efficient mapping of applications to MPSoCs is a highly challenging task. This tutorial focuses on the application domain of wireless communication, including radio standards like UMTS, LTE, and beyond. MPSoCs for wireless have evolved to complex systems, providing homogeneous and heterogeneous processing elements next to multi-level memory hierarchies. This mix of hardware facilities is needed to obtain the required processing speed and power efficiency, but makes the system hard to implement from the application programmer perspective. With a blend of experienced academic and industrial speakers, the tutorial discusses the major challenges and state-of-the-art solution approaches in embedded MPSoC programming. It covers DSP oriented programming models and languages, spatial/temporal task mapping and scheduling techniques, as well as optimized parallel code generation. Further key aspects are SDR based modem implementation, handling of real-time constraints, advanced code instrumentation for code analysis, vector processing, and utilization of MPSoC communication architectures. The presentations will be rounded up by concrete tool descriptions and/or demos, showing how many of the presented concepts can assist embedded MPSoC programmers already today, and what is needed for the future.

B Manufacturing, Design, and Test of 2.5D- and 3D-Stacked ICs

Organiser: Erik Jan Marinissen, IMEC, BE

Speakers: Armin Klumpp, Fraunhofer EMFT, DE
Paul Franzon, North Carolina State U, US
Erik Jan Marinissen, IMEC, Belgium

At the intersection of advanced semiconductor processing and advanced packaging, the semiconductor industry is preparing itself for vertical interconnection of multiple stacked dies by means of through-silicon vias (TSVs). TSVs are conducting nails which extend out of the back-side of a thinned-down die and enable the vertical interconnection to another die. TSVs are high-density, low-capacity interconnects compared to traditional wire-bonds, and hence allow for many more interconnections between stacked dies, while operating at higher speeds and consuming less power. TSVs promise to revolutionize the semiconductor industry by enabling the creation of new generations of 'super chips', in both "2.5D" (active dies placed on and interconnected by a passive silicon interposer containing TSVs) and "3D" (towers of vertically stacked dies, interconnected by TSVs).

Recently, several leading semiconductor companies, foundries, assembly houses, and their suppliers, have done product announcements, as public testimonials that 2.5D- and 3D-Stacked ICs are real. Consequently, process, design, and test engineers throughout the entire semiconductor industry are preparing for this product reality.

This comprehensive tutorial consists of three parts: (1) manufacturing, (2) design, and (3) test. In Part 1, on manufacturing, the tutorial presents the process steps related to TSV manufacturing, as well as the subsequent steps of wafer thinning, back-side processing, bonding, and packaging. In Part 2, on design, the tutorial focuses on the differences in architectures and design implementation between the conventional (2D) and the new 2.5D- and 3D-chips, and the status of the design automation in this field. We illustrate the challenges and benefits of stacked design by means of several case studies. In Part 3, on test, the tutorial presents various 3D test flows and the cost modelling thereof and new defects and test generation for stacked ICs. We also cover the challenges and emerging solutions with respect to test access, in terms of wafer probing and on-die design-for-test hardware.

C Mixed Signal IC Design and Test: Challenges, Solutions and Industry Practice

Organiser: Gordon Roberts, McGill U, CA

Speaker: Gordon Roberts, McGill U, CA

This tutorial will describe the challenges and practices of analog/mixed-signal (MS) design and test at a level that is suited to the non-expert. In much the same way that a computer program is written to implement a specific signal-processing algorithm, e.g., remove noise from a camera image, the goal of any analog/mixed-signal circuit design is to do a similar operation using a specialized transistor implementation. In order to appreciate the details of analog/MS design, this tutorial will begin with a review of the MOS transistor and how its voltage biasing and aspect ratio control its performance attributes. These insights will serve as the guide to the design of several basic circuits such as current sources and amplifiers. During the manufacturing, such circuits experience large variations in behavior. To solve these problems, unit-ratio component design together with feedback methods are used.

While negative feedback has largely been responsible for the success of analog/MS circuits, it is also the reason that these circuits are so difficult to test. As analog/MS circuits today achieve incredibly high levels of performance, it does so by operating at the limits of what component matching and negative feedback can provide. This, in turn, makes testing these circuits extremely time consuming, i.e., costly, as the test information lies with very small signals buried in noise. This tutorial will describe both production test techniques and several DFT approaches used in practice today; as well we shall look ahead into possible DFT approaches being discussed today in various research circles.

D1 **Multi-Core Platforms for Mixed-Critical Embedded Systems**

Organiser: Prof. Dr.-Ing. Rolf Ernst, TU Braunschweig, DE

Speakers: Prof. Dr.-Ing. Rolf Ernst, TU Braunschweig, DE
Glenn Farrall, Infineon, UK
Jonas Diemer, TU Braunschweig, DE
Henrik Theiling, Sysgo, DE
Matthieu Lemerre, CEA, FR
Swapnil Gandhi, Delphi, DE
Madeleine Faugère, Thales, FR

Mixed-critical systems integrate safety-critical and non-critical applications on the same platform leading to conflicts in verification, certification and maintenance. Most of the more complex embedded systems e.g. in automotive, avionics, or industrial automation are evolving as mixed critical systems to meet non-functional requirements (cost, space, weight, reliability, ...). Multi-core platforms offer improved isolation opportunities over single-core processors, but many challenges must be addressed, such as in core-to-core communication, isolation of shared resources and error protection.

The tutorial presenters are from hardware and real-time operating systems providers, covering different industrial segments and approaches. They will give an overview on the challenges and on the state of the art in industrial and academic solutions and give an outlook on the possible use in larger many-core systems and other open challenges. Specifically, the presenters will discuss hardware mechanisms for virtualization, isolation, and core-to-core communication, and present implementation examples. On the software side, the tutorial will detail how these mechanisms are used to provide a safe environment for applications following domain specific standards like ARINC and AUTOSAR. PikeOS and PHAROS will be covered as examples. Finally, the tutorial will discuss key applications of the presented platforms, focusing on domain specific requirements in both design and certification for automotive, avionics and industrial automation domains. The presenters collaborate in the RECOMP project (Reduced Certification Costs Using Trusted Multi-core Platforms, http://www.recomp-project.eu), a large European ARTEMIS project that focuses on multi-core platforms for mixed-critical systems.

E1 **HW - SW Design and Verification for Safety Critical Electronic Systems: Theory, Normative and Industry Practices**

Organiser: Riccardo Mariani, YOGITECH SpA, IT

Speakers: Riccardo Mariani, YOGITECH SpA, IT
Carsten Gebauer, Robert Bosch GmbH, DE
Karl Greb, Texas Instruments, US
Pete Harrod, ARM Ltd, UK
Martin Schwarz, TTTech Computertechnik AG, AU
Andreas Buchwieser, Wind River Gmbh, DE
Ioannis Sourdis, Chalmers U of Tech., SE

Nowadays, many HW and SW components are used in safety-critical domains such as automotive, industrial, medical, railway and aerospace. It is expected than within 2020 a great majority of electronic systems will have to cope - directly or indirectly - with safety requirements. From HW point of view, the new technologies are bringing new fault models and failure modes; from SW point of view, the coexistence in the same HW device of several applications and

tasks with different safety requirements brings complex problems in terms of data and timing interferences. To handle this complexity, functional safety standards like IEC 61508 and ISO 26262 as other standards like DO-254 are giving requirements, evaluation metrics and methodologies to define "how much safe" shall be a given component or system.

The organizer will introduce the tutorial giving an overview of the theory behind functional safety and a practical overview of the functional safety standards like IEC 61508, ISO 26262 and DO-254/DO-178B. The presenters from industries will give the audience a complete view on how HW and SW safety critical electronic systems are conceived and designed: from a system perspective, from a silicon vendor perspective, from a SW perspective, from an on-chip and off-chip safety network perspective and from a processor perspective. From research point of view, the tutorial will include a presentation about most advanced researches ongoing in the safety-critical domain, on behalf of the DeSyRe FP7-ICT project consortium.

F1 On Variability and Reliability; Dynamic Margining and Low Power

Organisers: Fadi J Kurdahi, UC Irvine, US
Ahmed M. Eltawil, UC Irvine, US
Amin Khajeh, Qualcomm Inc, US

Speakers: Fadi J Kurdahi, UC Irvine, US
Ahmed M. Eltawil, UC Irvine, US
Amin Khajeh, Qualcomm Inc, US

The design for manufacturing and yield (DFM&Y) is fast becoming an indispensable consideration in today's SoCs. Most current flows only consider manufacturability and yield at the lowest levels: process, layout and circuit. As such, these metrics are treated as an afterthought. With advanced process nodes, it has become increasingly expensive - and soon prohibitive - to guarantee bit level error free chips. The challenge now is to design reliable systems using chips that may have some faults. This has lead to approaches that consider DFM&Y at the system level where more benefit can be reaped, and to consider the problem across the design layers. This tutorial covers cross layer approach to design for DFM&Y spanning from the application all the way to manufacturing, overviews various techniques being explored today, and demonstrates its effectiveness on key applications including wireless communication systems (using WCDMA as the transmission physical layer), and multimedia applications (H.264). In addition, we explore the viability of cross-layer DFM&Y at the architectural level, focusing on processor caches. Experimental results that confirm the viability of such an approach will be presented and discussed. The results confirm that there is a significant opportunity for cross-layer error exploitation, resulting in an expanded design space with interesting design points that would otherwise have not been discovered or considered by SoC designers. The tutorial then proceeds to describe a scalable, unified statistical model that accurately reflects the impact of random hardware failures (embedded memory as an example) due to power management policies on the overall performance of a communication system. This enables system designers to efficiently and accurately determine the effectiveness of novel power management techniques and algorithms that are designed to manage both hardware failure and communication channel noise, without the added cost of lengthy system simulations that are inherently limited and suffer from lack of scalability.

G1 Demystifying Board-Level Test and Diagnosis

Organiser: Krishnendu Chakrabarty, Duke U, US

Speakers: William Eklow, Cisco Systems Inc, US
Krishnendu Chakrabarty, Duke U, US

The gap between working silicon and a working board/system is becoming more significant and problematic as technology scales and complexity grows. The result of this increasing gap is failures at the board and system level that cannot be duplicated at the component level. These failures are most often referred to as "NTFs" (No Trouble Founds). The result of these NTFs can range from higher manufacturing costs and inventories to failure to get the product out of the door. The problem will only get worse as technology scales and will be compounded as new packaging techniques (SiP, SoC, 3D) extend and expand Moore's law. This is a problem that must be solved, yet, little effort has been applied up to this point. This tutorial will provide a detailed background on the nature of this problem and will provide DFT and test solutions at both the component and board/system level.

This tutorial is part of the annual IEEE Computer Society TTTC Test Technology Educational Program (TTEP) 2012

D2 New Challenges and Technologies Behind Cloud Computing

Organisers: Marcello Coppola, STMicroelectronics, FR
Miltos Grammatikakis, TEI of Crete, GR

Speakers: Yiannis Kompatsiaris, ITI-CERTH, GR
Jan Kiszka, Siemens, DE
Marcello Coppola, ST Microelectronics, FR
Bernard Candaele, Thales, FR
Huy-Nam Nguyen, BULL, FR

The continuous decline of energy sources together with the rapid growth of commercial applications, such as online transaction processing, scientific computing, social media mining, multimedia-oriented web-services, and search engines, provided to end-users from a wide range of devices challenge existing large-scale computing, communication and data storage infrastructures. In fact, the extremely high energy costs of data centers due to vast server grids, power supply and cooling infrastructures, as well as the level of carbon emission in IT and communication technologies continue to raise public awareness, forcing policy makers to prepare legislation incentives towards green computing.

Existing high end multicore server architectures, such as IBM's Power7, Oracle Sun PowerNap and AMD's "Magny Cours" Opteron processor incarnation, already benefit from powerful out-of-order cores combined with large on-chip cache hierarchies and/or rapid, non-intrusive, automated power (or thermal) model transitions between a high-performance active state and a near-zero-power idle state.

Within this context, this tutorial examines challenges within the vast and ever-changing world of cloud technology, focusing on architectural trends, state-of-the-art design methodology and tools, full virtualization solutions, real-life applications and evolution towards large-scale, "green" multi-server cloud infrastructures. Key technology must be carefully weighed against potential costs, flexibility, scalability, performance, carbon footprint, power-efficiency, security, fault tolerance and overall quality of experience in accessing server content.

This tutorial brings together well-established actors from leading-edge companies, universities and research centers, firmly rooted in business realities and in tune with future research trends and evolution. The combined strongly-connected presentations will foster well-focused information exchange to the audience, providing ample space for questions.

E2 The Device-to-System Spectrum – A Tutorial on IC Design with Nanomaterials

Organisers: Deming Chen, U of Illinois at Urbana-Champaign, US
Subhasish Mitra, Stanford U, US

Speakers: Deming Chen, U of Illinois at Urbana-Champaign, US
Subhasish Mitra, Stanford U, US;
Eric Pop, U of Illinois at Urbana-Champaign, US
Naresh Shanbhag, U of Illinois at Urbana-Champaign, US

Nanomaterials such as carbon nanotubes (CNT), graphene nanoribbons (GNR), nanowires (NW), and other emerging electronic elements, have the potential to revolutionize nanoelectronics by enabling favorable device properties, novel functionality, or ultra-low power operation. Nanomaterials have a significant potential for building superior devices, routing structures, and interconnects. These nanomaterials may be essential to sustaining the advancement of electronic systems or bringing in new functionality. Yet they are also facing unique challenges, such as higher defect rate and nanomaterial-specific process-related variations. Such unique challenges are more apparent for scaled CMOS technologies as well. Given the great promise of nanotechnolgy for the development of future electronics and the general interest of research on this fast-growing area, we gathered a group of experts to give a tutorial to share the vision, present the promises and challenges, inspire the audience, and enable further development of nanotechnology.

In this tutorial, we will first present the current state-of-the-art fabrication, modeling and operation of several new nanoscale device and circuit components. These will include design and fabrication of GNR transistors, CNT transistors, CNT logic, low-power phase-change memory, etc. While these innovations are fundamental for the establishment of nanotechnology, the true impact for electronic systems demands that we translate these device and component-level capabilities into system-level benefits. Therefore, the second focus of this tutorial strives to bridge the gap between nanoelectronic device research and nanosystems design. These include nanosystems prototype design and modeling, statistical design approaches, and error-resilient designs targeting high performance, high reliability, and low power. The ultimate goal is to expose the potential of nanoelectronic technology and the unique challenges it presents across the whole device-to-system spectrum. Tutorial presentations will encompass four segments, with the first two segments focusing on nanodevices and nanocomponents, and the last two segments on nanosystems and system-level optimizations on performance, power and reliability.

F2 Design Methodology and Techniques in Production Low-Power SOC Designs

Organiser: Kaijian Shi, Cadence Design Systems, US

Speakers: Kaijian Shi, Cadence Design Systems, US
Thomas Buechner, IBM, DE

Power has become a critical metric and key differentiator in sub-65nm SOC designs, due to growing power density driven by technology scaling and chip integration. This tutorial provides an overview of the low-power design methodologies and techniques in production SOC design perspective, emphasizing on the real design considerations and impact on chip success. We shall discuss pros and cons of the methods and techniques considering impacts on chip design schedule, yield, and overall power-performance target. We shall also discuss about design guidance and recommendations in various design steps and decision making points, based on our years of successful experience in production low-power SOC designs.

This tutorial is organized in two parts. In the first part, we shall overview power related challenges in sub-60nm SOC design and state-of-the-art techniques to reduce chip power. We shall give a holistic view from chip level to system and application levels. Practical industrial examples will be used to show how power savings can be achieved in modern SOC, processors and computer systems. In the second part, we shall describe production low-power design methodology and techniques particularly the power-gating and the voltage & frequency scaling which are the two advanced power reduction methods used effectively in sub-65nm production low-power designs. We shall explain when, where and how these methods and techniques are applied to a chip according to the design goals and time-to-market requirement. We shall also overview production low-power design methodology and flow with power intent descriptions.

G2 Testing Embedded Memories in the Nano-Era: Fault Models, Tests, Industrial Results and BIST

Organisers Said Hamdioui, TU Delft, NL
Ad J Van de Goor, TU Delft and ComTex, NL

Speakers Said Hamdioui, TU Delft, NL
Ad J Van de Goor, TU Delft and ComTex, NL

The cost of memory testing increases with every generation of new memory chips. New technologies are introducing new defect mechanisms that were unknown in the past. Precise fault modeling to design efficient tests is therefore essential in order to keep the test cost and test time within economically acceptable limits, while keeping a high product quality.

The objective is to provide attendees with an overview of fault modeling, test design, BIST and BISR for memory devices in the nano-era. Traditional fault modeling and recent development in fault models for current and future technologies are covered. Systematic methods are presented for designing and optimizing tests, supported by industrial results from different companies (e.g. Intel, ST, Infineon) and for different technology nodes (e.g. 0.13um, 65nm). Impact of algorithmic (e.g. data-background) and non-algorithmic (e.g. voltage) stresses is explored in order to get better insight in the test effectiveness. State-of-the art and novel BIST architectures are covered; special attention is given to the optimization of address generator designs as they typically consume considerable BIST area overhead. BISR and redundancy analysis are also discussed. Finally, future challenges in memory testing are highlighted.

This tutorial is part of the annual IEEE Computer Society TTTC Test Technology Educational Program (TTEP) 2012