

TUTORIALS

A System-Level Modelling, Analysis and Synthesis of Embedded Multi-Core Designs

Organiser: Daniel D Gajski, CECS, UC Irvine, US
Speakers: Andreas Gerstlauer, UT Austin, US
Juergen Teich, University of Erlangen-Nuremberg, DE
Wolfgang Ecker, Infineon Technologies, DE
Samar Abdi, CECS, UC Irvine, US
Christian Haubelt, University of Erlangen-Nuremberg, DE
Mike Meredith, Forte Design Systems, US
Martin Speitel, Fraunhofer IIS, DE
Daniel D Gajski, CECS, UC Irvine, US

The continuous increase in size, complexity and heterogeneity of embedded system design has introduced new challenges in their modeling and implementation. Multi-core embedded system design requires high speed models for early validation and performance evaluation. As a result, electronic system level (ESL) modeling has moved up in abstraction from cycle accurate RTL to timed and untimed transaction-level models (TLMs), which are typically based on C or C++. However, the open question is how to get from a high level C based system description to a SW and HW implementation? The goal of this tutorial is to answer such questions and to provide system designers and managers with new insight into ESL modeling concepts and synthesis techniques so that they can drastically increase their productivity in the future.

In this tutorial, we will cover the key concepts and state of the art tools for C Based system design using industrial case studies such as MP3 player, JPEG encoder and digital radio designs. We will discuss TLM semantics for automatic model generation, methods for automatic design space exploration, and HW/SW synthesis. We will also present new advances in commercial C to RTL techniques, and their integration within complete ESL modeling, exploration, and synthesis frameworks.

This tutorial is targeted towards embedded SW and HW developers, engineers who use or are interested in using ESL design tools, managers of system designers, and verification engineers.

B Energy Harvesting Systems: Principles, Modelling and Performance Optimisation

Organisers: Tom Kazmierski and Steve Beeby, University of Southampton, UK
Speakers: Christoph Grimm, Vienna Technical University, AT
Thomas Herndl, Infineon Technologies, AT
Tom Kazmierski, University of Southampton, UK
John Parker, Perpetuum Ltd, UK
Paul Mitcheson, Imperial College, UK
Cees Links, GreenPeak Technologies, UK

This tutorial provides an introduction to operating principles of kinetic micro-generators and associated electronics, modelling and power optimisation of mixed-technology energy harvester systems using modern hardware description languages. The fundamental analysis will be supported by real world industrial case studies of energy harvester systems in sensor applications that highlight the challenges involved. The focus of this tutorial will be on kinetic energy harvesting, which converts movement or vibrations into electrical energy and is the subject of a considerable research interest.

Implications of harvested power in autonomous electronic systems design will be explained as well as general background necessary in the modelling of mixed-technology electrical/non-electrical systems. Extensive coverage will be given to the design of both the power conditioning electronics and the autonomous system electronics including operating strategies and wireless communications. Hardware description language descriptions will be used to illustrate the principles of mixed-technology system modelling and optimisation where electrical, mechanical and magnetic parts interact. The tutorial will also outline design techniques to obtain optimal power conditioning circuits that maximise available energy and strategies that minimise power consumption and enable operation.

The main objective of this tutorial is to provide the electronic designer with the knowledge required to fully exploit the energy harvesting revolution.

C Correct-by-Construction Embedded Software Synthesis: Formal Frameworks, Methodologies, and Tools

Organisers: Jean-Pierre Talpin, INRIA Rennes-Bretagne-Atlantique, FR
Sandeep K Shukla, Virginia Tech, US and INRIA Rennes-Bretagne-Atlantique, FR

Speakers: Stephen Edwards, Columbia University, US
Hiren Patel, University of California at Berkeley, US
Denis Laroudie, Geensys, FR
Sandeep Shukla, Virginia Tech, US and INRIA Rennes-Bretagne-Atlantique, FR
Jean-Pierre Talpin, INRIA Rennes-Bretagne-Atlantique, FR

Embedded Systems are ubiquitous. In applications ranging from control systems in avionics, automotive, and industrial process to handheld PDAs, cell phones, and bio-medical prosthetic devices, you can find embedded computing devices. According to some, embedded computing has seen some of the most explosive growths among all the fields in computing today. Some of the applications running on such embedded computing platforms are also safety-critical, real-time, and require absolute guarantee of correctness, timeliness, and dependability. As a result, such applications must be designed with utmost care, verified for functionality, satisfy real-time constraints, and must be properly endowed with reliability/dependability properties. These requirements pose new challenges to system designers. Much work has been done over the last few decades towards achieving the goal of sound programming models, abstractions, and underlying system architecture supporting these models. In this tutorial we will introduce the audience with three major approaches to specification driven embedded software synthesis/construction: synchronous programming based approaches, models of computation based approaches, and an approach based on concurrent programming with a co-design focused language.

The synchronous programming approaches include Esterel, Lustre and SIGNAL and various other programming models based on synchrony assumption. These approaches provide a multi-step refinement based strategy from concurrent programming abstraction to concrete implementation while preserving correctness all the way. The model of computations (MoC) based approach is exemplified in the Ptolemy framework where heterogeneous embedded software specification is captured with interacting domains specialised for various kinds of computational and communication primitives. The other approach is exemplified in SHIM, a concurrent programming model with asynchronous communication primitives meant to capture specification of an embedded system to be implemented using a hardware/software co-design approach. While all the developments in these three approaches will be presented and compared to make the audience understand the benefits and applicability of these approaches, an industrial environment along with language and semantics for automotive electronics hardware/software system design will be presented finally to provide the audience with an idea on how these academic approaches are applied in the industry.

The intended audience includes industrial engineers, and researchers; as well as academic researchers interested in safety-critical embedded software design in a systematic way with reduced verification burden, and guarantees of correctness. Fundamental understanding on embedded systems and its challenges will be beneficial but are not prerequisite to understand this tutorial

D1 New Developments and Trends in Networks on Chip

Organiser: Marcello Coppola, STMicroelectronics, FR
Marcello Lajolo, NEC, US
Riccardo Localelli, STMicroelectronics, FR
Partha Kundu, INTEL, US
Pascal Vivet, CEA-LETI, FR

Networks on Chip (NoCs) are rapidly becoming the mainstream architectural paradigm to cope with backend issues and application requirements. People attending this tutorial will discover the latest industrial developments in NoC design and short and long term research directions in SoC for consumer applications and multi-core/many-core systems. This tutorial includes a short history on on-chip interconnects followed by three Industry presentations (ST, INTEL, NEC) and one Academic presentation (CEA-LETI). We cover NoC design fundamentals, industry need and why the time is right for NoC design.

The tutorial is intended for engineers, managers, students, academics designers and managers who are interested to learn more about support technologies (SoC and massively parallel many-core architectures, communication-centric design, globally asynchronous and locally synchronous logic, programming languages and paradigms) that are necessary for realising the full potentials of NoC design.

E1 Low Power Design under Parameter Variations

Organiser: Swarup Bhunia, Case Western Reserve U, US
Speakers: Swarup Bhunia, Case Western Reserve U, US
Kanak B Agarwal, Austin Research Lab, IBM, US
Kaushik Roy, Purdue University, US

Low power design under parameter variations has emerged as a major design challenge in the nanometre regime. Design considerations for low power operation and variation tolerance typically impose contradictory requirements. Low-power design techniques such as voltage scaling, dual- V_{th} assignment and gate sizing can have large negative impact on parametric yield under process variations. Moreover, temporal parameter variations caused by temperature fluctuations or device degradations can adversely affect the reliability of low power designs.

This tutorial aims at providing comprehensive coverage on modelling, analysis and circuit/architecture level design methodology for low power and variation-tolerant logic circuits, memory and systems. It will address temperature-aware design; dynamic adaptation to temperature variations and device degradations such as Negative Bias Temperature Instability (NBTI). Post-silicon calibration and healing techniques using voltage scaling, adaptive body biasing or clock stretching will also be discussed. Current industry practices along with industrial data on process variability, variability characterisation, modelling and mitigation techniques will be presented.

This tutorial is targeted towards practicing VLSI design engineers, technical managers, tool developers, researchers and students working in the area of low power and variation tolerant VLSI design.

F1 Cross-Layer Approaches to Designing Reliable Systems using Unreliable Components

Organiser: Fadi Kurdahi, University of California at Irvine, US
Speakers: Fadi Kurdahi, University of California at Irvine, US
Nikil Dutt, University of California at Irvine, US
Ahmed Eltawil, University of California at Irvine, US
Sani Nassif, IBM Austin Research Labs, US

Design for manufacturing and yield (DFM&Y) is rapidly becoming an indispensable consideration in today's SoCs. Most current flows only consider manufacturability and yield at the lowest abstraction levels: process, layout and circuit. As such, these metrics are treated as an afterthought and drive only local design changes. With advanced process nodes the impact of design details on performance increases, making it increasingly expensive -- and soon prohibitive -- to guarantee 100% error free chips. The challenge now is how to design reliable systems using circuits that may have faults due to manufacturing process fluctuations, exasperated by environmental factors such as voltage and temperature variations. This leads to approaches that consider DFM&Y at the system level where more benefit can be reaped, and to consider the problems and solutions across the design layers. This tutorial covers a cross layer approach to design for DFM&Y spanning from the application all the way to manufacturing, overviews the various techniques being explored currently, and demonstrates its effectiveness on key applications including wireless, multimedia and imaging.

G1 Advanced Testing and Test Driven Self-Tuning of Mixed-Signal/RF Circuits and Systems

Organisers: Dimitris Gizopoulos, University of Piraeus, GR
Kaushik Roy, Purdue University, US
Speakers: Jacob A. Abraham, University of Texas at Austin, US
Abhijit Chatterjee, Georgia Institute of Technology, US

The impact of process uncertainties on high-speed mixed-signal/RF circuits arising from technology scaling must be taken into account during circuit design and manufacturing test. This tutorial develops advanced mixed-signal/RF test techniques that allow complex specifications to be evaluated at low cost. Low test cost is achieved through the use of low cost (low speed) external test instrumentation, the use of "intelligent" tester load boards and the use of on-chip circuitry to facilitate built-in test (BIT). The feasibility of completely autonomous built-in self-test is discussed and design tradeoffs are presented. The tradeoffs involved in using conventional specification test methods vs. advanced alternative test techniques based on test learning are analysed. It is shown how such advanced built-in self-test (BIST) techniques can be used to perform post-manufacture self-tuning of mixed-signal/high-speed/RF circuits for the purpose of achieving high yield under large manufacturing process variations. Performance/power constrained self-tuning techniques are discussed and it is shown how tuning can be performed for complex specifications (such as EVM) for which conventional built-in test driven tuning is impractical from tuning time considerations.

The intended audience includes mixed-signal/RF design and test engineers, test practitioners and managers, students and academics.

This tutorial is part of the IEEE Computer Society TTTC Test Technology Educational Program (TTEP) 2009

D2 Formal and Semi-formal Methods for Correctness and Robustness

Organiser: Görschwin Fey, University of Bremen, Bremen, DE
Speakers: Dominique Borrione, TIMA Laboratory, Grenoble, FR
Rolf Drechsler, University of Bremen, Bremen, Germany
Emmanuelle Encrenaz-Tiphene, Université Paris VI, Paris, FR
Goerschwin Fey, University of Bremen, Bremen, DE

Formal methods can guarantee 100% coverage which cannot be achieved by simulation based verification. The tutorial presents an overview of formal methods and their use along the design process. Underlying reasoning techniques are studied in detail. Then, the tutorial focuses on three recent innovative developments:

- (i) Dynamic verification of functional specifications
Assertions are written in formal languages: PSL and SVA are IEEE standards. The tutorial will focus on the tools that compile properties into synthesizable modules. This approach automates the debugging using both simulation and emulation, and still can be founded on formal techniques.
- (ii) Combined verification of functional and timing specifications
Until now functional correctness and timing requirements are either verified in a separate way or by simulation. The tutorial shows how formal methods verify both aspects within a single timed model.
- (iii) Robustness analysis
Due to continuously shrinking feature sizes fault tolerance will be required for almost any electronic device in the near future. The final part of the tutorial focuses on techniques to assess the robustness of a design. A fully automatic formal approach for robustness checking is introduced.

The purpose of the tutorial is an introduction to formal methods and to show their ability to solve recent and newly emerging design challenges.

E2 3D Integration – Opportunities, Challenges and Industry Readiness - Perspectives from Design, Manufacturing and EDA

Organiser: Charles Chiang, Synopsys Inc., US
Speakers: Louis Liu, TSMC, TW
Pol Marchal, IMEC, BE
Riko Radojcic, Qualcomm, US
Subarna Sinha, Synopsys Inc., US

Today's SoCs/SiPs face numerous design challenges as increased integration of system components on a single die stretches the limits of technology and design capacity. 3D integration, where multiple dies are stacked and interconnected in the vertical dimension using through-silicon vias (TSVs), is probably the best hope for carrying ICs along (and even beyond) the path of Moore's Law in the 21st century. However, successful adoption of 3D ICs by the semiconductor industry will require a killer application as well as modifications to design methodologies, EDA tools and manufacturing flows to enable 3D IC design. This tutorial will present the perspectives of industry leaders in design, manufacturing and EDA on the above-mentioned challenges and opportunities presented by 3D ICs.

Topics covered in this tutorial will include the architectural opportunities gained by designing in 3D, the technical challenges associated with the design and fabrication of such circuits, the "3D readiness" of the design, manufacturing and EDA sectors as well as the steps being undertaken by them towards commercial large scale use of 3D ICs. An overview of process technology aspects of 3D integration will be presented. Successful design of 3D ICs will require numerous changes to the current EDA tools as well as development of new tools. Concrete examples of new tools - such as technology tuning tools to co-optimize process technology and chip design requirements to define and validate the design enablement rules and models and architectural-level tools for co-exploring technology specifications and design options to define the optimum chip architecture for a given 3D process - will be described. A summary of the changes needed in key EDA stages as well as novel algorithms/methodologies for some of them will be presented. These include techniques for thermal modeling and algorithms for thermal-aware 3D floorplanning and a novel parasitic extraction methodology for 3D ICs. In addition, TCAD data for test structures for noise and reliability will also be presented.

The intended audience for this tutorial is EDA engineers who are interested to make their commercial tools “3D ready”, researchers looking for interesting relevant topics for research as well as designers looking to migrate their own designs into 3D.

F2 Power Optimised Design Techniques for Modern FPGAs

Organisers: Juanjo Noguera, Xilinx Inc., IE
Juergen Becker, University of Karlsruhe (TH), DE
Speakers: Juergen Becker, University of Karlsruhe (TH), DE
Robert Esser, Xilinx Inc., US
Michael Huebner, University of Karlsruhe (TH), DE
Juanjo Noguera, Xilinx Inc., IE

Power minimisation is becoming a major design goal for systems implemented using Field Programmable Gate Arrays (FPGAs). The two main contributions to FPGA power consumption are: due to the intrinsic FPGA micro-architecture and silicon implementation, and how designers implement applications on FPGAs including the design techniques applied and the particular CAD tools used. This tutorial addresses the second category, i.e. those aspects a designer can influence and covers power-aware best design practices targeting currently available modern FPGAs.

The tutorial covers the basics of why FPGA power consumption is important and explains the sources of power consumption in modern FPGAs. The focus of this tutorial is to describe power optimised FPGA design techniques at multiple levels: (1) circuit and architecture design techniques; (2) back-end CAD tools (e.g., synthesis and place & route); and (3) front-end CAD tools (e.g., HW/SW partitioning and partial reconfiguration). Finally, the tutorial introduces domain-specific languages as a research topic to facilitate the programmability of FPGAs and describe how FPGAs can achieve a competitive power advantage when compared to other programmable platforms (e.g., DSP and NPU) in different application domains (e.g., video and networking domains). The presented power-aware FPGA design techniques are demonstrated using practical examples and experimental results.

The tutorial is addressed to hardware and system engineers as well as to researchers interested in FPGA power optimization.

G2 Reliability, Availability and Serviceability of Networks-on-Chip

Organisers: Dimitris Gizopoulos, University of Piraeus, GR
Kaushik Roy, Purdue University, US
Speakers: Erika Cota, UFRGS, Porto Alegre, BR
Marcelo Lubaszewski, UFRGS, Porto Alegre, BR

This tutorial presents an overview of the issues related to the test, diagnosis and fault-tolerance of NoC-based systems. First, the characteristics of the NoC design (topologies, structures, routers, wrappers, and protocols) are presented, as well as a summary of the terms used in the field and an overview of the existing industrial and academic NoCs. Then, the challenges to test, diagnose and tolerate faults in NoC-based systems are discussed. Current test strategies are presented: re-use of the network for core testing, test scheduling for the NoC reuse, test access methods and interface, efficient reuse of the network, and power-aware and thermal-aware NoC-based SoC testing. In addition, the challenges and solutions for the NoC (interconnects, routers, and network interface) test and diagnosis are presented. Finally, since quality-of-service is one of the main challenges for the NoC use, fault-tolerance techniques for the NoC are discussed.

The intended audience are professionals, students, design and test engineers, and researchers interested in having introductory and intermediate knowledge on recent advances in test, diagnosis and fault-tolerance of integrated systems based on NoCs.

This tutorial is part of the IEEE Computer Society TTTC Test Technology Educational Program (TTEP) 2009