A Low Power SoC Design: Best Practice – and what's next?

Organiser: David Flynn, ARM, UK

Speakers: David Flynn, ARM, UK
Alan Gibbons, Synopsys, UK
David Jacquet, ST-Ericsson, FR
José Pineda de Gyvez, NXP, NL
Praveen Raghavan, IMEC, BE

Low power design is an ever growing challenge and concern for designers, especially in battery powered consumer devices where both active and standby energy consumption and usable lifetime impact strongly on product success.

Active power management approaches such as Dynamic Voltage and Frequency Scaling, DVFS, are well understood but tough to put into practice industrially. Static power mitigation techniques such as Power Gating are increasingly required on advanced technology nodes with significant sub-threshold and gate leakage currents.

Over the last four or so years Electronic Design Automation tools have standardised on significant support for "multi-voltage" design and implementation and "power-intent" extensions to annotate RTL designs for supply voltage, power switching, isolating or clamping interfaces and inferring state retention.

In order to exploit these low power techniques effectively a system-level or architectural perspective increasingly matters in order to control and optimise the hardware layers of power management.

Understanding basic industrial best practice now, and how more advanced techniques can build on the existing multi-voltage tools support is a goal of this DATE 2011 tutorial that brings together presenters with a wealth of experience and expertise in real-world low power system design.

The invited presentations will cover theory and practice, starting with building blocks and components at the IP level and outlying methodologies for design, implementation and verification with the evolving power intent format standards (CPF and UPF).

Worked examples from multi-media and wireless application space will be addressed with a strong focus on practical design that works in production, and the areas of software and hardware partitioning and optimisation.

The tutorial will round off with a look at what is next on the low power roadmap for energy optimal design and research approaches to technology scaling and resilient and error tolerant design.

B Electronic System Level Design and Verification

Organisers: Thomas Bollaert, Mentor Graphics, US
Carole Dunn, Mentor Graphics, US

Speakers: Thomas Bollaert, Mentor Graphics, US
Yvan Desmartin, STMicroelectronics, FR
Michael Fingeroff, Mentor Graphics, US
Bernhard Niemann,
Fraunhofer Institute for Integrated Circuits, DE

Today's advanced designs have grown too massive and complex to cost-effectively design and verify using traditional RTL methods alone. Electronic System Level (ESL) design methods give designers the answers they need to overcome this dilemma. And of all the ESL

methodologies flooding the market, the one technology that has gained considerable traction and use is high-level synthesis (HLS).

This tutorial explores the practical application of HLS with presentations from knowledgeable HLS users. In addition, Michael Fingeroff, the author of the High-Level Synthesis Blue Book, shares his guidance for how novices can go to experts by simply following the best practice coding examples shown in the book. Join us for a full-day of exploring how the promise of HLS is becoming a reality.

C Manufacturing, CAD and Thermal-Aware Design for 3D System-on-Chip Design

Organisers: David Atienza, EPF Lausanne, CH
Yuan Xie, Penn State U, US

Speakers: David Atienza, EPFL, CH
Tanay Karnik, Intel, US
Jeonghee Shin, IBM, US
Sachin S Sapatnekar, U of Minessota, US
P Leduc, CEA-LETI, FR
Yuan Xie, Penn State U, US

Three-dimensional integration (3D) with stacked chips is emerging as an attractive solution for overcoming even further the barriers to interconnect scaling, thereby offering an opportunity to continue performance improvements using CMOS technology, with smaller form factor, higher integration density, and the support for the realisation of mixed-technology chips in multi-processor system-on-chip (MPSoC) designs. Consequently, 3D Integration technologies have become the focus of the recent semiconductor and IC design R&D activities world-wide, and recent advances in process technology have brought 3D technology to the point where it is feasible and practical, and it has raised widespread interest in the semiconductor and fab-less industry. However, adding the third dimension implies the redesign of EDA tools and design/architectural techniques to fully explore new approaches and address the additional challenges of the 3D manufacturing complexity. Moreover, 3D stacking creates even higher power and heat density, leading to degraded performance if thermal management is not handled properly in novel system-level design flows, which can combine mature multi-tier IC stacking manufacturing, appropriate 3D computer-aided design (CAD) tools and sensible hardware/software thermal-aware architectures for 3D MPSoCs.

The goal of this full-day tutorial is to provide a thorough overview of the complete design flow for 3D MPSoCs and mixed-design ICs, namely, starting from the 3D integration process technology up to the application and market drivers for 3D products, while covering the different system-level EDA related challenges in 3D MPSoCs, architectures for 3D MPSoC design and, finally, thermal modeling and management of MPSoC at the system level.

Thus, this tutorial brings together leading 3D IC and MPSoC design experts in both industry/research centers (Intel, IBM, CEA-LETI) and academia (PennState, Univ. Minesota and EPFL), from both US and Europe, in order to cover in a comprehensive and structured way the key aspects of 3D technologies, design and products.

D1 MPSoC Hardware/Software Architectural and Design Challenges/Solutions

Organiser: Bernard Candaele, Thales, FR

Speakers: Geert Vanmeerbeeck, IMEC, BE
Kari Tiensyrjä and Jari Kreku, VTT, FI
Axel Jantsch, KTH, SE
Dimitrios Soudris, ICCS, GR
Bernard Candaele and Sylvain Aguirre, Thales, FR

Multi-core SoCs are rapidly becoming main stream but raise numerous architectural hardware and software challenges in designing and programming. The tutorial will review state of the art, research and developments and next challenges for HW/SW architectures and design space exploration to implement multi-cores/many cores SoC or FPGA platforms for wireless and multimedia applications.

Mapping software onto multi-processor platforms requires efficient parallel programming techniques whilst achieving non-functional requirements. The fundamentals, design steps and alternative programming models to implement such embedded applications onto multi-cores are discussed and presented.

The need to accommodate a large number of applications on these massively parallel computing platforms requires the system engineer to quickly evaluate the performances of application mappings. The tutorial will review mainstream evaluation techniques based on simulation, abstract workload and processing capacity models.

On-chip and in-package memory organisation and efficient data management are key to high performance. The tutorial will review various memory architectures and techniques to address space management, cache coherency, memory consistency, and dynamic application specific memory allocation techniques.

The tutorial will conclude with a case study in telecom, reviewing current related studies and next challenges.

In summary, the tutorial will focus on the following topics:

- High level design space exploration of heterogeneous and customisable multi-core architecture based on abstract architecture and workload models representing the applications;
- Options and implications of the on-chip memory architecture covering central, distributed, private, and shared organisations;
- On-chip memory architectures and application specific dynamic memory allocation techniques for the considered multi-core architectures.

E1 On-chip Interconnect for New Generation of SoC

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

Speakers: Reinig Helmut, Infineon, DE
Luca Carloni, Columbia University, US
Raj Yavatkar, Intel, US
Michael Dimelow, ARM, UK
Philippe D'audigier, STMicroelectronics, FR

On-chip interconnect design for embedded systems is moving away from traditional shared buses, such as AMBA AHB, IBM CoreConnect and STMicroelectronics STBus, towards sophisticated point-to-point communication architectures enriched by a set of services. The proliferation of multicore system-on-chip (SoC) and the desire to effectively consolidate applications in standardised platforms are driving forces of on-chip interconnect evolution. On-chip interconnects can benefit from knowledge transfer of PC,

communication server and IT industry by embracing several well-known technologies. However, due to critical challenges from complex applications and deep submicron technology, we must completely rethink our use cases and provide further enhancements for on-chip interconnects. New principles, constraints and processes are necessary for optimising system-level reuse, wire density, gate complexity, reliability, performance, scalability, power-efficiency and thermal management within the on-chip domain.

Within this context, this tutorial aims to evaluate the effectiveness of emerging innovative design methodologies, tools and leading-edge industrial practices for the evolution of on-chip interconnects in current state-of-the-art and next generation multicore SoC. Active researchers and well-established design managers firmly rooted in business realities and in tune with future trends will present important aspects and challenges within the vast and ever-changing world of on-chip interconnect technology, providing real life SoC examples.

A balanced, well-organised structure with strongly-connected presentations accompanied by numerous illustrations and easy-to-understand examples will foster well-focused information exchange to the general embedded community, providing ample time for questions. The target audience includes mainly CS/EE professionals (e.g. hardware engineers, SoC architects and managers), as well as semiconductor vendors and investors.

F1 Renewable Energy: Solar Power Generation, Conversion and Delivery

Organisers: Roberto Zafalon, STMicroelectronics, IT
Francesco G Gennaro, STMicroelectronics, IT

Speakers: Francesco G Gennaro, STMicroelectronics, IT
Gianluca Gigliucci, Enel, IT
Salvatore Lombardo, IMM - CNR, IT
Lars Bomholt, Synopsys, CH

Efficient generation, distribution and management of electric energy can reduce consumption in Europe from 20% to 30% by 2020 and CO2 emission in the same order of magnitude to achieve the Kyoto protocol targets and to limit the energy cost increase.

Solar energy generation proves to be a sustainable and efficient policy for global CO2 reduction. When coupled with the "Smart Energy Grid", it gives a reliable, flexible, accessible and yet cost-effective power source.

The Tutorial focuses on the solar energy supply chain, ranging from solar cells to innovative energy harvesting techniques, high efficiency power conversion and energy delivery through the smart grid.

The state-of-the-art and future trends of solar cells technology will be discussed, from standard Si cells to ultra-thin (20um) Si wafer PV cells, Si heterojunction cells, novel architectures (e.g. back-contact), new materials (for ARC, and passivation dielectrics), screen printing, laser processing and fully printable Dye-Sensitized-Solar Cells (DSSC). Next to the technology overview, the current status and features of system simulation of photovoltaic solutions using Synopsys TCAD Sentaurus for solar cell design and Saber for PV system simulation will be provided, suitable to model the thermally induced variations in performance.

Power conversion and control system design will be then addressed for a variety of cell technologies, with a special focus on Maximum Power Point Tracking operation.

Finally, energy delivery and transport will be the next challenge in the Smart Grid, with regards to management and control of electricity's exchange by means of communication infrastructure,

G1 Power-Aware Testing and Test Strategies for Low Power Devices

Organisers: Yiorgos Makris, Yale U, US
Dimitris Gizopoulos, Piraeus U, GR

Speakers: Patrick Girard, LIRMM/CNRS, FR
Nicola Nicolici, McMaster U, CA
Xiaoqing Wen, Kyushu Institute of Technology, JP

Managing the power consumption of circuits and systems is now considered as one of the most important challenges for the semiconductor industry. Elaborate power management strategies, such as voltage scaling, clock gating or power gating techniques, are used today to control the power dissipation during functional operation. The usage of these strategies has various implications on manufacturing test, and power-aware test is therefore increasingly becoming a major consideration during design-for-test and test preparation for low power devices. This tutorial provides knowledge in this area. It is organised into three main parts. The first one gives necessary background and discusses issues arising from excessive power dissipation during test application. The second part provides comprehensive knowledge of structural and algorithmic solutions that can be used to alleviate such problems. The last part surveys low power design techniques and shows how these low power devices can be tested safely without affecting yield and reliability. EDA solutions for considering power during test and design-for-test are also discussed in the last part of the tutorial.

This tutorial is part of the annual IEEE Computer Society TTTC Test Technology Educational Program (TTEP) 2011

D2 Model-based MPSoC Architecture Synthesis for Highly-demanding Embedded Applications

Organiser: Lech Jozwiak, TU Eindhoven, NL

Speakers: Lech Jozwiak, TU Eindhoven, NL
Menno Lindwer, SiliconHive, NL
Jan Madsen, TU Denmark, DK

The recent spectacular progress in modern nanoelectronic technology enabled implementation of very complex multiprocessor systems on single chips (MPSoCs) and created a big stimulus towards development of high-performance systems for various highly-demanding embedded applications. In result, the increasingly complex and sophisticated MPSoCs are required to perform real-time computations to extremely tight schedules, with high demands regarding energy, power, area, and cost efficiency. Moreover, they are required to be flexible enough to enable reuse among different product versions, adherence to evolving standards or user requirements, and easy modification during their development or even their field use. This all results in serious design and development challenges, such as: multi-objective MPSoC optimisation, resolution of numerous complex design tradeoffs, reduction of the design productivity gap, time-to market and development costs without compromising quality, etc. The opportunities created can effectively be exploited only through use of more adequate application-specific system architectures and more integrated system IP modules, supported by new system-level design methods and EDA tools. This tutorial focuses on mastering the automatic architecture synthesis and application mapping for heterogeneous customisable multi-processor systems-on-chip (MPSoCs) based on configurable and extensible application-specific instruction-set processors (ASIPs). The MPSoC design technology based on adaptable ASIPs is able to deliver high performance, high

flexibility and low energy consumption at the same time. It is relevant for a very broad range of modern applications and applicable to several implementation technologies. The tutorial presents the results of our analysis of the main problems that have to be solved and challenges to be faced in design of such heterogeneous customisable MPSoCs for modern demanding applications. It discusses the problems of architecture synthesis and application mapping involving multi-objective optimisation, adequate exploitation of multiple trade-offs, and coherent development of computing, communication and memory sub-systems for complex hard real-time embedded MPSoCs, as well as, proposes the model-based semi-automatic architecture synthesis methods and EDA-tools that enable effective and efficient solution of these problems.

E2 Design and Verification Challenges for Automotive Electronics

Organiser: Riccardo Groppo, Centro Ricerche FIAT, IT

Speakers: Riccardo Groppo, Centro Ricerche FIAT, IT
Patrick Leteinturier, Infineon, DE
Erwan Hemon, Freescale, FR
Pascal Caunegre, Freescale, FR
Manfred Thanner, Freescale, FR
Franco Toto, STMicroelectronics, IT
Davide Appello, STMicroelectronics, IT
Matteo Sonza Reorda, Politecnico di Torino, IT

The increasing complexity of automotive electronics, both in terms of electronic control systems' architecture and the new generation of complex ICs, must meet ever more demanding vehicle's requirements. On one hand the request for higher computational performances is due to more complex real-time mathematical models meant to achieve a better system control (i.e. improved combustion process for CO2 reduction); furthermore these mathematical models can also be designed to serve as "virtual sensors" in order to limit an overall system cost which is increasing day by day due to fuel economy and clean transportation increasing demands: this has been the driving factor of multi-core microcontrollers' architecture. The request for a better system partitioning and computational balance has also lead to the development of a new generation of mixed A/D intelligent ICs. On the other hand the need for improved reliability and, most of all, safety levels together with the stringent target cost, typical of the automotive business, is creating one of the biggest challenges at the silicon makers. In fact, automotive systems must have a high level of functional safety and fault tolerance built-in to ensure dependable and predictable operation over the lifetime of the product, which will be more than 10 years. The automotive-specific ISO 26262 safety standard defines the maximum acceptable failure rates and fault tolerant features that should be considered in the design and implemented in the development process. As a consequence modern design techniques and methodologies must be used in order to achieve those targets and predict system behavior in case of fault. Hence the tutorial will present basic and more advanced topics that cover various aspects of these challenges in a comprehensive fashion. It will start with the OEM view of the challenges, then advanced design methodologies, high level modeling, verification and statistical screening techniques will be described. The approach towards functional safety will require adequate efforts in terms of formal verification. Finally, an exhaustive testing automation and verification phase must be planned.

F2 Overcoming CMOS Reliability Challenges: From Devices to Circuits and Systems

Organiser: Yu (Kevin) Cao, Arizona State U, US
Subhasish Mitra, Stanford U, US

Speakers: Yu (Kevin) Cao, Arizona State U, US
Georges Gielen, KU Leuven, BE
Subhasish Mitra, Stanford U, US
Sani Nassif, IBM, US

With extreme miniaturisation of CMOS circuits, factors such as transient errors, device degradation, and variability induced by manufacturing and operating conditions are becoming important. While design margins are being squeezed to achieve high energy efficiency, expanded design margins are required to cope with process variability and device degradation. Even if error rates stay constant on a per-bit basis, total chip-level error rates grow with the scale of integration. Moreover, difficulties with traditional burn-in can leave early-life failures unscreened. This tutorial will present a wide range of topics that cover various aspects of these challenges in a comprehensive fashion: starting from device-level reliability modeling and characterisation, all the way to techniques for designing globally-optimised robust circuits (both digital and analog) and systems. Innovations in EDA tools to overcome reliability challenges will also be discussed. By focusing on resilience techniques across multiple abstraction layers (often referred to as cross-layer resilience techniques), this tutorial will illustrate new opportunities in designing cost-effective robust systems of the future.

G2 Testing TSV-Based 3D Stacked Ics

Organisers: Yiorgos Makris, Yale University, US
Dimitris Gizopoulos, Piraeus U, GR

Speakers: Erik Jan Marinissen, IMEC, BE
Yervant Zorian, Synopsys, US

Three-dimensional stacking of multiple integrated circuits has benefits in terms of combining heterogeneous technologies and achieving a small footprint. The semiconductor industry is preparing itself to make a major step forward in three-dimensional stacking, now that the technology of TSVs is becoming available. TSVs are conducting nails which extend out of the back-side of a thinned-down die, enabling the vertical interconnect to another die. TSVs are high-density, low-capacitance 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. TSV-based 3D technologies enable the creation of a new generation of 'super chips' by opening up new architectural opportunities. 3D-SICs combine a smaller form factor and lower overall manufacturing costs with many other compelling benefits, and hence their technology is quickly gaining ground.

Like all micro-electronic products, 3D-SICs need to be tested for manufacturing defects incurred during their many, high-precision, and hence defect-prone manufacturing steps. These tests should be both effective and cost-efficient. Solutions regarding test flow, test contents, and test access need to be developed before 3D-SICs can be brought to the market. Next to all basic and most advanced test technology issues, 3D-SICs have some unique new test challenges of their own. These challenges include (1) development of new fault models and corresponding tests for TSV-based interconnects and new 3D-induced intra-die defects, (2) wafer probing on small and numerous micro-bumps and/or TSV tips and pads under stringent damage requirements, (3) handling of and probing on wafers with thinned-die stacks, (4) the design, partitioning, and optimisation of DfT architectures that span across multiple dies, and (5) optimisation of the test flow for maximum effectiveness and lowest cost.

This tutorial is part of the annual IEEE Computer Society TTTC Test Technology Educational Program (TTEP) 2011