Tutorials

A1 Sidestepping Performance Bottlenecks and Design Crises with Better Than Worst-Case Design
Organiser: Todd Austin, Michigan U, US
Speakers: Todd Austin, Michigan U, US
Valeria Bertacco, Michigan U, US
Krisztian Flautner, ARM Ltd, UK

This tutorial introduces the audience to state-of-the-art design methods that address the correctness and reliability challenges of deep-submicron silicon, targeting both circuit and architecture aspects of the design flow. The focus is on a new design strategy, called Better Than Worst-Case design, which couples complex design components with simple reliable checker mechanisms. By delegating the responsibility of correctness and reliability to the checker, it becomes possible to quickly build designs that are provably correct and that address effectively performance and reliability constraints. Two exemplary better than worst-case designs are presented. In addition, a complementary design technique, called typical-case optimization (TCO), is introduced as a way to take advantage of the relaxed design constraints on the core component. Finally, the tutorial highlights some of the new opportunities and challenges posed to CAD tools by this new design methodology.

The tutorial is intended for designers and CAD engineers interested in novel techniques for robust design, their potential applications, and their implications to system design and CAD tools.

B1 Substrate Noise Coupling: A Potential Barrier to Complete Systems-on-Chip
Organisers: Piet Wambacq and Stephane Donnay, IMEC, BE
Speakers: Marcel Pelgrom, Philips Research, NL
Adil Koukab, EPFL, Lausanne, CH
Jajeet Roychowdhury, Minnesota U, Minneapolis, US
Geert van der Plas, IMEC, BE

The different circuits of a system on a silicon chip share the same semiconducting silicon substrate. Via this substrate, parasitic couplings can arise which disturb sensitive circuits. A typical example is substrate coupling from switching digital circuits to sensitive analogue circuits. This problem is not straightforward to model and to suppress, such that it can cause extra design iterations, over-dimensioning or complete failures.

Substrate noise coupling has three aspects: first, generation of noise. On a mixed-signal IC, this is typically caused by the switching of the digital circuits. Another example is the generation of substrate noise by analogue circuits that handle large signals such as on-chip power amplifiers. To compute the amount of generated noise, macromodelling approaches are required. Second, there is the aspect of propagation through the substrate from the generating circuit(s) to the sensitive circuits that are affected. To generate a circuit-level model, which can be simplified if needed, different approaches exist, some of them being available in tools. Finally, there is the aspect of impact on analog and RF circuits. These circuits can be impacted via different entry points such as the bulk of the different transistors, the ground and supply lines, passive components in the circuit.

To introduce the substrate noise coupling problem, this tutorial first covers some case studies on real-life ICs that suffer from the noise coupling problem. Next, state-of-the-art approaches are discussed on the analysis of the three aspects of substrate noise coupling mentioned above. Both research approaches are discussed as well as industrial experience and experiments to suppress substrate noise coupling.

C1 The POSIX Profile for Embedded Real-Time Operating Systems
Organiser: Michael Gonzalez Harbour, Cantabria U, ES
Speakers: Michael Gonzalez Harbour, Cantabria U, ES
Mario Aldea Rivas, Cantabria U, ES

The POSIX standard defines services for large operating systems, including the support for real-time requirements. The full standard would require a large powerful platform, inappropriate for small embedded systems. For that reason, standard subsets or profiles have been defined, including one for small embedded systems.

The tutorial will discuss the main real-time operating system services defined in the POSIX.13 minimum real-time profile. These services allow application developers to write portable applications that meet their real-time requirements, and that may be implemented on small embedded systems.

The tutorial will provide detailed examples that will be demonstrated using MaRTE OS, a free software implementation of the POSIX minimum real-time system profile.

D1 Geometric Programming and its Applications to EDA Problems
Organiser: Stephen Boyd, Stanford U, US
Speakers: Seung Jean Kim, Stanford U, US
Sunderarajan Mohan, Barcelona Design, US

Over the past 15 years there has been a revolution in convex optimization, spurred by new solution methods that efficiently solve even large scale problems, and the discovery of a wide range of new applications. The tutorial will begin with an overview of convex optimisation and the new solution methods, and will then focus on geometric programming (GP), a more specific family of optimisation problems, which has recently been used for a variety of EDA problems, including the design of analogue and mixed circuits ranging from op-amps to whole PLLs and ADCs, and digital circuit sizing. We will cover the basics of GP, including some standard tricks used to transform problems to GP form, and a powerful extension called generalised geometric programming (GGP).

The focus will be on GP modelling, the task of approximately formulating a design problem in GP format. We will cover related topics such as fitting empirical data or functions in a form compatible with GP, and how to handle discrete constraints on some of the variables.

We will then look at a variety of problems in analogue and digital circuit design to which GP and GGP can be applied. Examples will include amplifier design problems, including formulation of bias equations, small signal, and other constraints, as well as design over corners, and joint electrical and physical design and digital circuit sizing problems, ranging from simple gate scaling and Elmore delay-based device and wire sizing to joint optimisation of gate scaling and supply and threshold voltages, and heuristics for statistical optimization.

E1 Memory Test Challenges — A Practical and Implementation View of BIST and other DFT Techniques
Organiser: Dimitris Gizopoulos, Piraeus U, GR
Speakers: Vinay Jayaram, Texas Instruments, US
Sherry Lai, Texas Instruments, US
Theo Powell, Texas Instruments, US

The task of optimally configuring memory BIST and implementing it is becoming increasingly difficult in today's deep sub micron designs. Designers are facing issues with floorplanning constraints, synthesis and optimisation flows, timing closure, integration with the design structure and overall execution, all coupled with reduced cycle-times.

The purpose of this tutorial is to educate the audience with the challenges associated with implementing memory BIST on actual designs. Topics would include memory BIST approaches, configuration techniques, design guidelines, automation capabilities, synthesis and timing closure tips, physical design considerations and overall implementation and execution pointers. In addition, the tutorial will also cover debugging approaches that can be used while testing devices on ATEs and some topics on memory diagnosis and repair techniques.

This tutorial will be based on the experiences of the authors in implementing at-speed memory BIST on many deep sub-micron ASIC designs.

This Tutorial is part of the IEEE Computer Society TTTC Test Education Program – TTEP 2005

A2 FPGA-based DSP Implementation versus Traditional DSPs
Organiser: Michael Schwarz, PLC2, Freiburg, DE
Speakers: Eugen Krassin, PLC2, Freiburg, DE

FPGAs have seen many developments since their first introduction, but the same fundamental structure remains. Although the inclusion of memory as a function was not present until early nineties, it became such a useful feature in DSP processing that it has become a standard. The FPGA is a fully SRAM based device, which is configured following application of power. The building blocks can be configured to provide many functions, with programmable interconnect used to join these small functions to form larger functions and systems. Therefore the FPGA provides a huge degree of freedom for the creation of DSP processing functions.

The most common DSP function implemented in FPFAs is the Finite Impulse Response filter. This is a nice algorithm that help us to understand many of the DSP techniques that can be employed in FPGAs. We do need to understand what mathematical calculations need to be performed and understand the typical requirements of the algorithm in systems.

The structure of the FIR filter is well known, but some “extra knowledge” can be very useful when making implementation decisions. Stating the number of taps, sample rate and sample size will impact the DSP implementation technique used. This tutorial will illustrate and compare various techniques for implementation of DSP functions in FPGAs. The theoretical section will be accompanied by a “Real World Design” based on Xilinx FPGAs.

B2 Interconnect Design for the Nanometer Era
Organiser: Sachin Sapatnekar, Minnesota U, Mineapolis, US
Speakers: Sachin Sapatnekar, Minnesota U, Minneapolis, US
Prashant Saxena, Intel Corp, US

With circuit technologies entering the nanometer regime, CAD tools are faced with an array of new challenges due to growing system complexities, aggressive speed and power specifications, and new effects that arise with shrinking geometries. Interconnects play a dominant role in determining the performance of high-performance digital systems, and it is imperative to pay careful attention to their design to ensure that high-performance systems meet the stringent specifications imposed upon them.

In the nanometer domain, new considerations must be accounted for in interconnect design. We will begin by listing the numerous challenges that designers must face in sub-90 nm technologies, with extrapolations based on industrial data that go up to the 32 nm node. Next, we will discuss issues in modeling interconnects, and survey the challenges involved in three types of global interconnect networks in post-RTL design: signal nets, clock nets and supply nets. Of these, signal nets are typically relatively small in size, but large in number. The supply and clock nets, on the other hand, are used for global distribution to numerous sinks all over the chip, and may have a tremendously large number of nodes. This tutorial is intended to provide a comprehensive view of the daunting problems associated with interconnect design in nanometer technologies, and will present approaches that may be used to overcome these challenges.

C2 Automotive Electronics: From Embedded Software to Distributed Embedded System Design
Organiser: Chidamber Kulkarni, Xilinx Inc, US
Speakers: Chidamber Kulkarni, Xilinx Inc, US
M. Kashif Imam, Volkswagen, US

It is now accepted that majority of all innovations in a vehicle (car) will be driven by electronics and software. In the near future it is also expected that almost 30 percent of cost of a new vehicle will be due to electronic components. Thus automotive electronics has become a central aspect of the total automotive design process. On the other hand, computing power and available memory doubles every 18 months and the available on-chip communication bandwidth doubles every 12 months. Thus there are two important technology roadmaps that need to be addressed in a single design process.

This tutorial will present the existing state-of-the-art in electronic vehicle architectures as well as the embedded software development process. Based on the two technology roadmaps, we will identify key limitations in existing architectures and present potential solutions to enable future vehicle architecture(s). In particular, we will identify above said challenges in existing reference designs and show how they can be re-engineered to overcome the limitations. Finally we will provide research directions for advanced design methodologies and tools for automotive electronics.

D2 Decision Procedures and their Applications to EDA Problems

Organiser: Fabio Somenzi, Colorado U, Boulder, US
Speakers: Fabio Somenzi, Colorado U, Boulder, US
Andreas Kuehlmann, Cadence Berkeley Labs, US

Decision procedures for various logics find increasing use in several areas of electronic design automation including verification, testing, synthesis, simulation, and layout. These procedures decide propositional logic, quantified Boolean formulae, and quantifier-free fragments of first-order logic. The field has received impetus from significant advances in the basic search technology and from the use of non-clausal representations. These improvements, pioneered by propositional SAT solvers, are finding their way into decision procedures for more expressive logics, which are often based on propositional abstractions. This tutorial will introduce the logics, review some of their applications in different fields, and discuss the algorithms that have been developed for their efficient decision.

E2 When Digital Becomes Analogue — Interfaces in High-Speed Test

Organiser: Dimitris Gizopoulos, Piraeus U, GR
Speaker: Wolfgang Maichen, Teradyne Inc, US

Successful high-speed digital test requires understanding the behavior of all components involved (device, test equipment, interface) as well as the interactions between them. The boundaries between digital and analogue disappear. Connections have to be treated as transmission lines, with every little imperfection impacting system performance, deteriorating test accuracy, and reducing yield.

Maintaining signal integrity is paramount to achieve meaningful results. This requires good knowledge of signal path behaviour, proper choice of termination techniques, and good interface design. While the tester itself is usually given and not easily modified, large improvements are possible by careful design of the interface board. Practical methods to quantify the performance of a given design by measurements and to build models based on these measurements are necessary as well.

This tutorial gives an in-depth introduction into the concepts, tools and methods necessary to successfully test and characterise today’s high-speed digital circuits and systems.

This Tutorial is part of the IEEE Computer Society TTTC Test Education Program – TTEP 2005