

Master Courses

M1 Technology, Circuits and EDA Techniques for Leakage Power Control in CMOS Designs

Organiser: Massoud Pedram, Southern California U, Los Angeles, US

Speakers: Massoud Pedram, Southern California U, Los Angeles, US

Farzan Fallah, Fujitsu Labs of America, US

Siva Narendra, Intel Corp, US

In many new high performance designs, the active leakage component of power consumption is comparable to the switching component. Reports indicate that 40 percent or even higher percentage of the total power consumption is due to the active leakage of transistors. This percentage will increase with technology scaling unless effective techniques are used to bring leakage under control. This master class will focus on technology, circuit and design automation techniques to accomplish this goal.

The first part of the master class provides an overview of basic physics and technology scaling trends that have resulted in a significant increase in the sub-threshold, gate tunneling and junction tunneling leakage currents. Distinction between active and standby leakage current will be introduced. An in-depth description of adaptive Vdd/Vt and multiple Vdd/Vt/Tox techniques for leakage minimisation in the presence of process variations will be provided and relevant methodologies and models will be described. This part will also highlight the use of high permittivity gate dielectric, metal gate and novel device structures for controlling the three leakage current components. The second part of the master class describes a number of design optimisation techniques for controlling the leakage current, including power/ground gating and precomputation-based signal guarding. It will also present runtime mechanisms for leakage control including reverse and forward body bias control, transition to minimum leakage state, and state assignment. The third part covers leakage-aware cell library design, RT-level synthesis techniques, low-leakage memory and cache circuits and architectures, and design of leakage-tolerant high performance datapath circuits.

M2 Transaction Level Modelling with the New OSCI SystemC TLM Standard

Organiser: Alan Fitch Aynsley, Doulos Ltd, UK

Speakers: John Aynsley, Doulos Ltd, UK

Stuart Swan, Cadence Design Systems, US

Alain Clouard, STMicroelectronics, FR

Mike Meredith, Forte Design Systems, US

One of the major advantages of SystemC is the ability to model at the transaction level (TLM), where the use of simple function calls in communication modelling brings gains in both coding productivity and simulation speed. But the introduction of new modelling paradigms brings with it the need for new modelling standards, and the introduction of TLM interface standards has become one of the top priorities within the Open SystemC Initiative (OSCI) in recent years.

The TLM standard is now here and is being adopted. The OSCI TLM interface standard extends the practical value of the SystemC class library by providing a standard modelling kit for the construction of TLM interfaces, thus reducing the work needed to construct new interfaces and increasing the opportunities for interoperability. At the same time, the release of version 2.1 of the SystemC class library has added new features which extend the utility of SystemC for transaction level modelling.

This tutorial is structured as follows: Part 1 introduces and explains version 2.1 of the SystemC class library, and is appropriate as a primer for those with limited knowledge of SystemC as well as for those who wish to be appraised of the latest developments. Part 2 introduces transaction level modelling and the new OSCI TLM Standard. Part 3 gives practical examples from industrial speakers of how SystemC is used for transaction level modelling in real-world SoC designs.