2011 CANDE Workshop Preliminary Program

November 10, 2011, Double Tree Hotel, San Jose, CA, USA

8:00-8:25 **Breakfast**

8:25-8:30 **Opening Remarks**

8:30-9:30 **Session I: Sense and Sensibility**
Chair: Gi-Joon Nam, IBM, USA

Using Infrared Imaging to Improve IC Design
Sherief Reda, Brown University

Synthesizing Logical Computation on Stochastic Bit Streams for Sensing Applications
Marc Riedel, University of Minnesota, Twin Cities

9:30-10:00 **Special Session: Student Research Highlights**

10:00-10:30 **Coffee Break + Poster Session**

10:30-12:00 **Session II: GPU for EDA Computing**
Chair: Kyle Rupnow, Advanced Digital Sciences Center, Singapore

EDA Computing with OpenCL for GPU Architectures and Beyond
Rasit Onur Topaloglu, GLOBALFOUNDRIES, Inc.

Leveraging GPU Computing for VLSI CAD and Beyond
Peng Li, Texas A & M University

GPU Evolution - Simplifying the Step to Heterogeneous Computing
Stephen Jones, NVIDIA, Inc.

12:00-1:15 **Lunch and Lunch Keynote**
EDA or the art of trading off apples, peaches and oranges
(... and why many-core GPU will barely move the needle)
Patrick Groeneveld, Magma

1:15-3:00 **Session III: Panel Discussion -- EDA Roadmap: Why and How?**
Moderator: *Farinaz Koushanfar, Rice University*

Panelists:
Andrew Kahng, University of California, San Diego
Juan-Antonio Carballo, NetLogic Microsystems
Dave Noice, Cadence
Leon Stok, IBM

3:00-3:15 **Coffee Break**

3:15-4:15 **Keynote Session**
Many-core GPU Computing - Current Victories and Coming Battles
Wen-Mei Hwu, University of Illinois, Urbana-Champaign

4:15-4:30 **Coffee Break**

4:30-5:30 **Session IV: New FPGA Advances**
Chair: Guy Lemieux, University of British Columbia, Canada

System-level Trends in FPGA Devices
Mike Hutton, Altera, Inc.

Recent Advances in Xilinx FPGAs
Tim Tuan, Xilinx, Inc.

5:30-6:30 **Evening Reception** (sponsored by IEEE CEDA)

Technical Session Details

8:30-9:30 **Session I: Sense and Sensibility**

Title: Using Infrared Imaging to Improve IC Design
Sherief Reda, Brown University

Abstract: Approximate design models, computational limits, and runtime and process variabilities all contribute towards deviating the results of CAD tools from actual IC post-silicon characteristics. Infrared imaging is a powerful technique to quantify post-silicon characteristics under realistic loading conditions. In this talk, an overview of basic concepts in infrared imaging and instrumentation will be first presented, and then a number of IC design applications will be discussed and demonstrated, including: hot spot detection, power mapping, reliability characterization, and Trojan detection. Results from test chips will be presented and analyzed.

Sherief Reda is an assistant professor of Electrical Sciences and Computer Engineering at the School of Engineering, Brown University, where he heads the SCALE laboratory (http://scale.engin.brown.edu). Professor Reda received his Ph.D. degree in computer engineering from University of California, San Diego, and B.Sc. degree from AinShams University, Cairo. His research interests include thermal / power modeling and management techniques, variability modeling and yield optimization, and physical design. He has received four best paper nominations and two best paper awards at DATE'02 and ISLPED'10.

Title: Synthesizing Logical Computation on Stochastic Bit Streams for Sensing Applications
Marc Riedel, University of Minnesota, Twin Cities

Abstract: Most digital systems operate on a positional representation of data, such as binary radix. A positional representation is a compact way to encode signal values: in binary radix, 2^n distinct values can be represented with n bits. However, operating on it requires complex logic: in each operation such as addition or multiplication, the signal must be "decoded," with the higher order bits weighted more than the lower order bits. We advocate an alternative representation: random bit streams where the signal value is encoded by the probability of obtaining a one versus a zero. This representation is much less compact than binary radix. However, complex operations can be performed with very simple logic. For instance, multiplication can be performed with a single AND gate. Also, because the representation is uniform, with all bits weighted equally, it is highly tolerant of soft errors (i.e., bit flips). In this talk, we will discuss a general method for synthesizing digital circuitry that computes on such stochastic bit streams, targeting sensing applications. We show that our method produces very simple hardware that is highly tolerant of soft errors, so suitable for very light-weight sensors operating in harsh environments.

Marc Riedel has been an Assistant Professor of Electrical and Computer Engineering at the University of Minnesota since 2006. He is also a member of the Graduate Faculty in Biomedical Informatics and Computational Biology. In 2004-2005, he was a lecturer in Computation and Neural Systems at Caltech. He has held positions at Marconi Canada, CAE Electronics, Toshiba, and Fujitsu Research Labs. He received his Ph.D. and his M.Sc. in Electrical Engineering at Caltech and his B.Eng. in Electrical Engineering with a Minor in Mathematics at McGill University. His Ph.D. dissertation titled "Cyclic Combinational Circuits" received the Charles H. Wilts Prize for the best doctoral research in Electrical Engineering at Caltech. His paper "The Synthesis of Cyclic Combinational Circuits" received the Best Paper Award at the Design Automation Conference. He is a recipient of the NSF CAREER Award.

10:30-12:00 **Session II: GPU for EDA Computing**

Title: EDA Computing with OpenCL for GPU Architectures and Beyond
Rasit Onur Topaloglu, GLOBALFOUNDRIES, Inc.

Abstract: Graphical processing unit (GPU) computing has been an interesting area of research in the last decade. While initial adapters of the technology have been from image processing domain due to difficulties in programming the GPUs, research on programming languages made it possible for people without the knowledge of low-level programming languages such as OpenGL develop code on GPUs. Two main GPU architectures from AMD (former ATI) and NVIDIA acquired grounds. While AMD adapted Stanford's Brook language and made it into an architecture-agnostic tool, NVIDIA brought CUDA framework to a wide audience. Until recently, it has not been possible to compile one GPU code across platforms. An opportunity came with idea of combining one or more CPUs and GPUs on the same die. Eliminating some of the interconnection bandwidth issues, this combination makes it possible to offload tasks with high parallelism to the GPU. The technological direction towards multicores for CPU-only architectures also requires a programming methodology change and act as a catalyst for better programming languages. Hence, a unified language that can be used both on multicore CPUs as well as GPUs gained interest. Open Computing Language (OpenCL), developed originally by the Khronos Group of Apple and supported by both AMD and NVIDIA, is seen as the programming language of choice for parallel programming. In this talk, I provide an introduction to EDA computing with OpenCL for current and future architectures with GPUs. I will end my talk by providing research directions in this area.

Dr. Rasit Onur Topaloglu received his B.S. degree with High Honors in Electrical and Electronic Engineering from Bogazici University in Turkey. He received his M.S. degree in Computer Science and Ph.D. degree in Computer Engineering from University of California at San Diego. Since 2005, he has been with Advanced Micro Devices and later on with GLOBALFOUNDRIES, which is an advanced semiconductor manufacturer with technology coverage down to 28 nm and research down to 14 nm. Dr. Topaloglu's research interests are VLSI design and CAD. He has over 40 international and refereed publications, four book chapters and an edited book, three granted and four pending patents, and a best paper award at IEEE International Symposium on Quality Electronic Design. He has chaired DAC Workshop on Parallel Algorithms, Programming, and Architectures (PAPA) in 2011. He serves in technical program committees for DAC and International Symposium on Physical Design. He serves as Semiconductor Research Corporation Science Area Coordinator for Computer Aided Design and Test as well as Circuits and Systems for GLOBALFOUNDRIES. He is also a Semiconductor Research Corporation Focus Center Research Program Associate for Concurrent Systems.

Title: Leveraging GPU Computing for VLSI CAD and Beyond
Peng Li, Texas A & M University

Abstract: The recent emergence of general-purpose GPU (graphics processing units) computing platforms offer appealing opportunities in leveraging these powerful SIMD (single instruction multiple data) machines for addressing a range of computational challenges. However, to be successful, one has to properly expose the "hidden" data parallelism in an application and develop "smart" algorithms to penetrate the SIMD barrier of the GPU architecture. A proper interplay between algorithm design and SIMD architecture consideration is essential for achieving good runtime performance. In this talk, we will show GPU-based algorithms can be developed to analyze large power delivery networks with excellent efficiency. We will also demonstrate the application of GPU computing to the simulation of large neuronal networks with biophysically plausible Hodgkin-Huxley type neuron models.

Peng Li received the Ph.D. degree in electrical and computer engineering from CMU in 2003, respectively. He is presently an associate professor at Department of Electrical and Computer Engineering, Texas A&M University. His research interests center on VLSI systems, design automation, parallel computing, and aspects of computational neuroscience and biology. His work has received several recognitions including an NSF CAREER Award and three IEEE/ACM Design Automation Conference (DAC) Best Paper Awards.

Title: GPU Evolution - Simplifying the Step to Heterogeneous Computing
Stephen Jones, NVIDIA, Inc.

Abstract: Heterogeneous computing - running a single program on two (or more) different computer architectures - offers the potential efficiency and performance of both worlds, but at the cost of increased program complexity. Ideally a program could achieve heterogeneous execution within a single, homogeneous programming model; in practice this requires a combination of software-aware hardware design, and heterogeneous-aware software design. In this talk, we will look at the co-evolution of GPU hardware and programming model design and where that trajectory will take us in the future - in particular with application to the large legacy code-bases typical in the EDA world.

Stephen Jones is an engineer in NVIDIA's high-performance computing group, working on developing the CUDA language for programming the GPU. His background is in computational fluid dynamics and plasma physics and he is a member of CUDA's parallel algorithms group, which develops the mathematical libraries which form part of the CUDA programming toolkit. His current work is focused on the evolution of the CUDA language to better support high performance computing applications, and the future GPU architectures required to enable it.

12:00-1:15 **Lunch and Lunch Keynote**

Title: EDA or the art of trading off apples, peaches and oranges (... and why many-core GPU will barely move the needle)
Patrick Groeneveld, Magma

Patrick Groeneveld is Chief Technologist at Magma Design Automation in San Jose, California. Since joining Magma in 1997, he has developed key parts of Magma's flagship physical synthesis products. Patrick received his MSc and PhD degrees from Delft university of technology (1987, 1991). Before starting at Magma, Patrick was associate professor at Delft University of Technology, where he was also associated with DIMES. From 2001 until 2005 he was full professor of Electrical Engineering at Eindhoven University. In his spare time, Patrick is the general Chair of the 2012 Design Automation Conference in San Francisco. Otherwise, Patrick enjoys being with his family, hiking, reading useless information and flying a Cessna 172.

3:15-4:15 **Keynote Session**

Title: Many-core GPU Computing - Current Victories and Coming Battles
Wen-Mei Hwu, University of Illinois, Urbana-Champaign

Wen-mei W. Hwu is the Walter J. ("Jerry") Sanders III-Advanced Micro Devices Endowed Chair in Electrical and Computer Engineering in the Coordinated Science Laboratory of the University of Illinois at Urbana-Champaign. Dr. Hwu received his Ph.D. degree in Computer Science from the University of California, Berkeley, 1987. His research interests are in the areas of architecture, implementation, software for high-performance computer systems, and parallel processing.

He is a Principal Investigator (PI) for the petascale Blue Waters system, is co-director of the Intel and Microsoft funded Universal Parallel Computing Research Center (UPCRC), and PI for the world's first NVIDIA CUDA Center of Excellence. At the Illinois Coordinated Science Lab, he is the director of the OpenIMPACT project, which has delivered new compiler and computer architecture technologies to the computer industry since 1987. He also serves as the Soft Systems Theme leader of the MARCO/DARPA Gigascale Silicon Research Center (GSRC) and on the Executive Committees of both the GSRC and the MARCO/DARPA Center for Circuit and System Solutions (C2S2). For his contributions to the areas of compiler optimization and computer architecture, he received the 1993 Eta Kappa Nu Outstanding Young Electrical Engineer Award, the 1994 Xerox Award for Faculty Research, the 1994 University Scholar Award of the University of Illinois, the 1997 Eta Kappa Nu Holmes MacDonald Outstanding Teaching Award, the 1998 ACM SigArch Maurice Wilkes Award, the 1999 ACM Grace Murray Hopper Award, the 2001 Tau Beta Pi Daniel C. Drucker Eminent Faculty Award, and the 2002 ComputerWorld Honors Archive Medal. He is a fellow of IEEE and of the ACM.

4:30-5:30 **Session IV: New FPGA Advances**

Title: System-level Trends in FPGA Devices
Mike Hutton, Altera, Inc.

Abstract: Recent years have seen many changes in FPGA technology, and one of the more pronounced ones is that FPGAs are no longer just glue logic but a central component of systems. This talk will describe recent solutions and promising new directions along with some remaining challenges. Specific topics will include newly announced embedded SOC devices, changes to FPGA CAD for system-level design (floorplanning, interconnect generation), new directions such as OpenCL design entry and commercial application of dynamic reconfiguration.

Mike Hutton is a Principal Design Engineer in the Office of the CTO at Altera Corporation. His research interests include FPGA architecture, CAD and high-performance applications on FPGAs. He is a member of the ACM SIGDA Technical Committee on FPGAs and Reconfigurable Computing, Associate Editor of IEEE Transactions on CAD, serves on the Program Committees of DAC, FPGA, FPL, FPT and other conferences and is past TPC-Chair for FPGA and SLIP. He has a Ph.D. from the University of Toronto, and has published 35 academic papers and holds more than 50 US Patents on FPGA and related technology.

Title: Recent Advances in Xilinx FPGAs
Tim Tuan, Xilinx, Inc.

FPGAs have evolved from simple, small arrays of logic blocks to massive, complex systems comprising billions of devices. With this fantastic growth comes a number of significant challenges, including manufacturability, power limitations, and design productivity. In 28nm process node, Xilinx is implementing several key technology innovations, such as 2.5D Stacked Silicon Interconnect, new power management solutions, and Extensible Processing Platforms. These technologies will help Xilinx FPGAs to cope with the challenges of ultra deep submicron technologies and to continue riding the waves of scaling.

Tim Tuan is a Sr. Staff Engineer in the CTO Office (aka Xilinx Research Labs) at Xilinx Inc. in San Jose. At Xilinx, he has worked extensively on analyzing and mitigating deep submicron challenges facing FPGAs using a wide range of techniques spanning process technology, circuits, architectures and CAD. Currently, he's working on high performance architectures and power management solutions for the next generation FPGAs. Prior to Xilinx, he was with Berkeley Wireless Research Center working on ultra low power wireless network processors. He has a MS in EECS from UC Berkeley and BS in Computer Engineering from University of Washington. He has published over 20 technical papers and 20 issued patents.