Guest Editorial System-Level Interconnect Prediction

THE IEEE-ACM International Workshop on System-Level Interconnect Prediction (SLIP—http://www.sliponline.org/) was first organized in 1999. Its central theme is to bridge the gap between detailed physical (material and geometric) interconnect properties, only known late in the design cycle, and *system-level* properties related to, e.g., power consumption, delay and yield, required for design space exploration and design optimization. This gap is usually bridged by a *predictive* approach: by developing a theoretical predictive model, by benchmarking and empirical model-construction, or by some combination of both. SLIP topics have touched interconnect analysis, model construction, and the application of models in an electronic design automation (EDA) setting as well as the evaluation and application of novel interconnect architectures and technologies.

Early on, the focus of SLIP was to a large extent on the prediction of logic-level ("sea of gates") layout results, often based on Rent's rule. More recently, the topics presented at SLIP have drifted along with the state-of-the-art EDA problems: both upwards and downwards in the design hierarchy. At the high end of the design hierarchy, the focus is on specific system-level interconnect architectures, such as network-on-chip (NoC) or bus architectures. At the lower end, we find the interconnect-related headaches of deep submicrometer design: signal integrity and process variation. Driven by the limitations of electrical interconnect, the quest for alternative interconnect technologies is also very active and reflected in recent SLIP programs. On an intermediate level, we find field-programmable gate array (FPGA) interconnect architectures—their design, their evaluation, and the prediction of design embedding results—as a constant throughout SLIP history.

For this Special Section, we have made a careful selection of seven regular papers and three brief papers. They reflect the four tracks described in the previous paragraph: system-level interconnect architectures, FPGA interconnect, prediction of circuit-level interconnect properties, and novel interconnect technologies.

This Special Section contains four papers on the design of interconnect architectures. Two of them focus on NoC. In "Polaris: A system-level roadmapping toolchain for on-chip interconnection networks," Soteriou *et al.* present a framework for rapidly exploring a wide range of NoC architecture alternatives when optimizing for a given application. The framework combines a set of tools developed by the authors. The speed of the exploration comes from the use of realistic synthetic traffic traces and simplified performance and cost projections. The second NoC paper, "Synthesis of predictable networks-on-chip-based interconnect architectures for chip multiprocessors," by Murali *et al.*, focuses on the importance of performance predictability when designing an NoC architecture for chip multiprocessors (CMPs). If congestion occurs in an NoC, communication delays can become unacceptably large. As the workload of CMPs is not known at design time, the authors guarantee predictable communication latency by ensuring that the network supports the maximally required throughput in the CMP system. Instead of greedily dimensioning all network links according to the overall peak throughput, the differences in individual link requirements are determined and accounted for to save network power consumption. Communication throughput is also addressed in "Design of an interconnect architecture and signaling technology for parallelism in communication" by Kim *et al.* However, focusing on off-chip communication, the motivation of this paper is to exploit the available resources to achieve maximal communication parallelism. To achieve this, the authors propose a new interconnect architecture and signalling technology, called source synchronous CDMA interconnect. It implements two simultaneous virtual channels on a single physical wire using three-level pulse amplitude modulation. The final paper on system-level interconnect solutions, "Energy/area/delay tradeoffs in the physical design of on-chip segmented bus architecture" by Guo *et al.*, addresses physical design exploration of application-specific communication architectures. In segmented buses, signal propagation is restricted to the segments required to reach the destination. In order to save power, heavily used segments are made as short as possible. The layout must, therefore, be driven with an activity-based cost, but without ignoring area and delay. Because it is hard to fix the balance between these three cost metrics in a generic way, the authors construct a pareto-optimal front by varying the contributions of each metric to the overall cost in an existing layout flow.

The three FPGA-related contributions to this special section have very diverse topics. "Post-placement interconnect entropy" by Feng *et al.* builds upon the SLIP tradition of Rent's rule-based analysis of interconnect requirements. The authors define the concept of interconnect entropy, representing an upper bound on the number of configuration bits required to implement a well-placed netlist which obeys Rent's rule. They use this metric to analyze how the required number of configuration bits scales for increasing design sizes and under which conditions it is possible to build a fully scalable FPGA cell architecture. "Predicting interconnect delay for physical synthesis in an FPGA CAD flow" by Manohararajah *et al.* also builds upon a more traditional SLIP topic: the inherent randomness of physical design and its implications on the achievable accuracy of predictions for individual wires. It turns out that, due to the specific architectural features of certain FPGAs, a portion of wire delay is determined by technology mapping, e.g., caused by which types of terminals a wire connects. As this information is available before physical design, the authors show how it can be usefully applied to better guide early physical synthesis. Finally, in "Routability prediction of network topologies in FPGAs" Saldaña *et al.* investigate the resource utilization when mapping different multiprocessor interconnection network topologies on

Digital Object Identifier 10.1109/TVLSI.2007.900756

FPGAs. In state-of-the-art FPGAs, a very large amount of interconnect resources are available anyway. The authors explore the question of how much connectivity you can provide under these conditions while still achieving the required number of processors per chip.

We have selected two papers treating the prediction of circuit-level interconnect properties. One cause of errors in digital systems is the degradation of power and ground voltages due to IR-drop and inductive phenomena, which are essentially stochastic phenomena caused by circuit activity and related to the geometry of the power/ground supply networks. To keep the impact of these effects within acceptable bounds, a careful design of power and ground distribution is required. The paper "Stochastic power/ground supply voltage prediction and optimization via analytical placement" by Kahng *et al.* proposes a fast, random-walk-based metric for stochastic power/ground (P/G) supply voltage degradation. The authors subsequently apply their metric to optimize P/G supply networks in an analytical placement framework. The second paper addresses another cause of variability: the presence of process variations which relatively have an increasing impact on interconnect delay. In "Fast variational interconnect delay and slew computation using quadratic models," Ye *et al.* present fast but accurate metrics for interconnect delay and slew analysis in the presence of process variability. Their delay and slew estimates for a given set of process parameters, are formulated as second-order polynomial deviations of the nominal values. The resulting models prove fast and sufficiently accurate for exploring interconnect-related delay variations.

One of the alternative interconnect approaches that has been under investigation for increasingly short-range interconnections is optical interconnect. The paper "Systematic simulation-based predictive synthesis of integrated optical interconnect," by O'Connor *et al.* addresses the design and evaluation of on-chip optical interconnects, the performance of which is determined by the effective codesign of electrical and optoelectronic components. It presents a systematic simulation-based synthesis framework for such interconnects, combining spice-level simulation models for all components. Using predictive component models, it can also be used to explore performance at future technology nodes or the overall performance impact of individual technological parameter improvements.

Overall, we have made a selection of high-quality innovative contributions, illustrating the broad spectrum of interconnect related design problems and predictive approaches to tackle them. The origins of these papers clearly reflect an active interest in interconnect-related design issues from both academia and industry. We would like to thank all contributing authors for submitting their work and the reviewers for helping to select the best of this issue's relatively high number of submissions.

JONI DAMBRE, *Guest Editor*
Electronics and Information Systems Department
Ghent University
Ghent, 9000 Belgium

MIKE HUTTON, *Guest Editor*
Altera Corporation
San Jose, CA 95134 USA

Joni Dambre (S'99–M'04) was born in Ghent, Belgium, in 1973. She received the M.S. degree in electronics engineering and the Ph.D. degree in computer engineering from Ghent University, Ghent, Belgium, in 1996 and 2003, respectively. Her dissertation treated Rent's rule-based prediction of VLSI wire length distributions and derived properties.

Since 2003, she has been a Postdoctoral Researcher with the Department of Electronics and Information Systems (ELIS), Ghent University. She has performed research on interconnect technology evaluation, with a focus on photonic interconnect, and on embedded system design automation. She teaches a course on circuit-level digital design. Her other research interests include platform and technology evaluation, EDA algorithms and their evaluation, neural computation as well as novel interconnection technologies.

Dr. Dambre is a member of the ACM. She has been on the SLIP program committee from SLIP 2004 onwards and was the technical program chair of SLIP 2006.

Mike Hutton received the Ph.D. degree in computer science from the University of Toronto, Toronto, ON, Canada, and the B.Math. and M.Math. degrees from the University of Waterloo, Waterloo, ON.

In 1997, he joined Altera Corporation, San Jose, CA, where he is currently a Principal Engineer. At Altera, he most recently worked on the architecture definition of the Stratix II, MAX II, and Cyclone II FPGA architectures and the HardCopy II structured application-specific integrated circuit (ASIC) architecture. He is the author of over 30 patents and 20 published papers in the areas of FPGA architecture and CAD. He serves on the program committees for FPGA, FPL, FPT, DATE, SLIP, IWLS, ISVLSI, and other conferences. His research interests include field-programmable gate array (FPGA) architecture and applications, algorithms for synthesis, timing analysis, and place&route, and graph-theoretic algorithms. He has been on the SLIP program committee from SLIP 2002 onwards. He was technical program chair of SLIP 2005 and general chair of SLIP 2006.