

Guest Editors' Introduction: Challenges for Reliable Design at the Nanoscale

R. Iris Bahar
Brown University

Mehdi B. Tahoori
Northeastern University

Sandeep K. Shukla
Virginia Tech

Fabrizio Lombardi
Northeastern University

■ **IT IS WITH GREAT PLEASURE** that we introduce this special issue on Advanced Technologies and Reliable Design for Nanotechnology Systems to the *IEEE Design & Test* readership. We have selected four articles to cover a wide spectrum of techniques and applications for the reliable design of nanoscale systems; the techniques aim to circumvent the high defect rates and transient errors expected in advanced nanoscale technologies. Written by outstanding researchers in the field, these articles cover experimental and speculative topics. As with all special issues, these topics only represent the techniques and methodologies available today.

Nanotechnology roadmap

To understand the timeliness of this special issue's topic area, we must look into the overall national research priority and the ongoing discussions in industry and academia about the future of computing, the longevity of Moore's law, and speculations about future technologies when Moore's law reaches its end.

One nanometer is a magical point on the dimension scale. Nanostructures are at the confluence of the smallest of human-made devices and the largest molecules of living things. Nanoscale science and engineering here refer to the fundamental understanding and resulting technological advances arising from the exploitation of new physical, chemical and biological properties of systems that are intermediate in size, between isolated atoms and molecules and bulk materials, where the transitional properties between the two limits can be controlled.¹

This quote from 2001 by Mihail Roco, a National Science Foundation senior advisor, shows not only that the technological and scientific inflection point—or

what former Intel CEO Andy Grove called “disruptive technology”—might be at our doorstep in the form of nanotechnology, but it also hints at the grand challenges that computer engineers and scientists face in the nanotechnology era. Robust computing systems design will soon meet quantum physical, probabilistic, and even biological phenomena. Guaranteeing reliability is more challenging than ever before, but the possible scale factor plays the role of a conflicting design force.

The single-electron scale of device design brings forth the challenge of probabilistic effects and uncertainty in guaranteeing zero-one-based computing. Minuscule devices imply billions of devices on a single chip, which might help mitigate the challenge of uncertainty by replication and redundancy, but will create a design and validation nightmare with the sheer scale.

Although much nanotechnology research is confined to material science, electrical engineering, quantum and device physics, chemistry, and so forth, computer engineers and scientists will soon see the effects that we've just described as nanostructured materials with unreliable and defective substrates begin to enter the mainstream of computer design. To build nanosystems using such unreliable nanoscale devices, computer engineers must overcome the following challenges:

- How do we guarantee reliability for a nanosystem that has many defects in the underlying substrate on which the logic and architecture was built?
- As past research has shown, too much redundancy can be counterproductive—the redundant parts will have the same uncertainties. How do we efficiently model and analyze such figures of merit as reliability, optimal redundancy level, and coding precision?
- Can we efficiently validate nanoscale designs in the

presence of uncertainty? This is important because some system functions might still be valid regardless of uncertainties in other parts of the system.

- How do we validate the functionality of billion-device nanosystems? In other words, how can current verification algorithms handle substantially larger designs?

Although these challenges are clear, the added layer of uncertainties and the sheer scale factor in nanosystems demand revolutionary breakthroughs in design technology. To raise awareness of the design and test challenges in nanoscopic-level design, we decided to bring forth this special issue. In this context, we must make a few more points regarding the imminent future of electronic design and design technologies. As the semiconductor industry continues to ride Moore's law, it seems that the end of the roadmap is near. Even though most manufacturers are preparing for 65-nm technologies as we publish this special issue, and 45-nm is not too far behind, we still hear technologists' growing concerns: the end of Moore's law, and even before the end of the roadmap, the challenges of defects, soft errors, and other transient faults. We summarize these concerns with a quote from 2002 by Mark Reed and James Tour:

... no one expects conventional silicon-based micro-electronics to continue following Moore's Law forever. At some point, chip-fabrication specialists will find it economically infeasible to continue scaling down microelectronics. As they pack more transistors onto a chip, phenomena such as stray signals on the chip, the need to dissipate the heat from many closely packed devices, and the difficulty of creating the devices in the first place will halt or severely slow progress.¹

This special issue

In the first article of this special issue, Thaker et al. discuss the conditions that allow for scaling fault tolerance in nanometer ranges. Starting from a mostly CMOS-based analysis, the article shows that, asymptotically, a novel arrangement called recursive triple modular redundancy (RTMR) does not necessarily yield the desired improvements in fault tolerance. For deep-submicron CMOS devices, the authors propose different noise models and show that RTMR can compensate for single events resulting from energetic particles. However, when device geometry substantially shrinks in size, RTMR circuits can be too large and offer only marginal improvements. Efficient implementation at such reduced sizes and accounting for other noise sources requires exploration of other tech-

niques. The authors discuss in detail various microarchitecture approaches that trade off reliability, speed, and area overhead for effective manufacturing.

Next, DeHon and Naeimi use a bottom-up assembly to consider defect rates for nanosystems, which are expected to be significantly higher than the defect rates currently encountered in the deep-submicron region. DeHon and Naeimi argue that this technology manufacturing must rely on modularity to provide sufficient flexibility and resource utilization in the likely occurrence of a highly defective fabrication process. Based on past experience in array-like architectures (such as FPGAs), the authors introduce a fine-grained programmable logic array architecture for implementing reconfiguration at different levels. Their proposed testing and defect-avoidance procedures for nanowires and programmable junctions within a 2D grid layout are amenable to nanotechnology implementation, including a molecular one. Using graphing and probabilistic methods, DeHon and Naeimi propose and analyze in detail different sparing allocation and matching/sparse-ness techniques for programmable resources and their interconnects. Moreover, the authors advocate on-chip test and reconfiguration support circuitry for efficient bootstrapping and designing to allow resources to handle a large number of defects.

In the third article, He, Jacome, and de Veciana introduce a reconfiguration fabric on a probabilistic basis to improve scalability in the fabrication process, a strict requirement at the nanometer scale because of the large number of devices and their likely failure. The authors propose a novel decomposition technique that relies on structural and behavioral criteria and abstractions; this technique instantiates yield with defect mapping by simultaneously satisfying other constraints, such as performance and chip cost. In particular, it preserves and enforces a very flexible hierarchy in the abstractions to enable the integration of fault tolerance and defect avoidance in design methodologies. This process is based on an "interface" that allows different components to integrate into a design phase. The authors also briefly introduce a speculative technique to address performance issues for reconfigurable systems.

In the last article, Han et al. briefly review von Neumann's classic N -tuple modular redundancy scheme, considering multiplexing and interwoven logic as additional criteria for the voting process. The authors use novel Markov chain models to analyze the arrangement by which multiplexing hardware is introduced and cascaded into stages. To assess noise sensitivity and

decreased fabrication tolerances, the authors characterize error correction features. Using these techniques, Han et al. show that an increased tolerance to defects is still possible in nanometer regimes. Moreover, they propose an interwoven redundant logic with randomized connectivity to fully account for the potential randomness of the very likely nanoscale interconnect in the assembly process. Through Markov chains and bifurcation analysis, the authors show that a comprehensive analytical platform is indeed possible by correlating the different features of these techniques. They study in detail a quadded logic arrangement with triple interwoven redundancy to fully assess the structure's viability at the nanoscale and establish a good agreement between simulation and analytical results.

WE SINCERELY HOPE that this special issue will be a reference publication for future research. These articles cover topics that are timely and important, and the authors have done an excellent job of presenting the material. We extend our sincere thanks to all the authors and reviewers. We also thank Rajesh Gupta, editor in chief of *IEEE Design & Test*, for allowing us to create this special issue. Finally, we give a special thanks to the editorial staff of the IEEE Computer Society for editing and assembling this issue. Please feel free to contact us if you have questions or comments. ■

■ Reference

1. M.A. Ratner and D. Ratner, *Nanotechnology: A Gentle Introduction to the Next Big Idea*, Prentice Hall PTR, 2002.



R. Iris Bahar is an associate professor in the Division of Engineering at Brown University. Her research interests include computer architecture; CAD for synthesis, verification, and low-power applications; and design and test issues for nanoscale systems. Bahar has a BS and an MS in computer engineering from the University of Illinois at Urbana-Champaign and a PhD in electrical and computer engineering from the University of Colorado at Boulder.



Mehdi B. Tahoori is an assistant professor in the Department of Electrical and Computer Engineering at Northeastern University. His research interests include nanotechnology, VLSI testing, CAD for VLSI, and fault-tolerant comput-

ing. Tahoori has a BS in computer engineering from Sharif University of Technology, Tehran, Iran; and an MS and a PhD in electrical engineering from Stanford University. He is the program chair of the 2005 IEEE International Workshop on Current- and Defect-Based Testing and cosponsor of the IEEE Test Technology Technical Council (TTTC) Technical Activity Committee on Testing of Nano-Based Devices and Systems.



Sandeep K. Shukla is an assistant professor of computer engineering and a College of Engineering Faculty Fellow at Virginia Tech, where he is also founder and deputy director of the Center for Embedded Systems for Critical Applications (CESCA) and director of the Formal Engineering Research with Models, Abstraction, and Transformations (Fermat) laboratory. His research interests include formal methods, hardware and software specification languages and methodologies, probabilistic system design and analysis, and reliability methods. Shukla has a BE in computer science and engineering from Jadavpur University, Calcutta, India; and an MS and PhD in computer science from State University of New York at Albany. He is an editorial board member of *IEEE Design & Test* and a senior member of the IEEE.



Fabrizio Lombardi holds the International Test Conference Endowed Professorship at Northeastern University, Boston. His research interests include the testing and design of digital systems, quantum and nanoscale computing, ATE systems, configurable and network computing, defect tolerance, and CAD for VLSI. Lombardi has a BSc in electronic engineering from the University of Essex, UK; and an MS in microwaves and modern optics, a Diploma in microwave engineering, and a PhD in engineering, all from University College London. He is the associate editor in chief of *IEEE Transactions on Computers*, an editorial board member of *IEEE Design & Test*, and the founding general chair of the IEEE Symposium on Network Computing and Applications. He also chairs the TTTC Technical Activity Committee on Testing of Nano-Based Devices and Systems.

■ Direct questions and comments about this special issue to Sandeep K. Shukla, Dept. of Electrical and Computer Eng., Virginia Tech, 302 Whittemore, Blacksburg, VA 24061; shukla@vt.edu.