

# Guest Editors' Introduction: Process Variation and Stochastic Design and Test

**T.M. Mak**  
Intel

**Sani Nassif**  
IBM Austin Research Laboratory

■ **PROCESS VARIATION** has been around since the advent of large-scale silicon processing, but has been kept under sufficient control to allow for the unprecedented 40 years of scaling predicted by Moore's law. However, it appears that the laws of physics are finally catching up with us.

As mainstream silicon manufacturing processes scale to and beyond the 65-nm node, we find that process variations are consuming an increasingly larger portion of design and test budgets. The manufacturing community still projects that CMOS will scale for another three to four generations, but the wavelength of the light we use to create devices is going to be an order of magnitude larger than the geometry we want to print. This necessitates the use of several resolution enhancement techniques, such as optimal proximity correction, phase-shifting masks, double exposure, and immersion optics. The use of these techniques will cause an ever-larger gap between the geometries drawn in CAD systems and the implemented devices, creating systematic differences between the two that lead to performance prediction uncertainty. Furthermore, as we scale the transistor further, physics-based random variations in the threshold voltage due to random dopant fluctuations, and in the channel length due to line edge roughness, add a random component of variability to the systematic layout versus silicon just mentioned. These variations play a significant part in sub-threshold leakage and other important device performance metrics. This uncertainty is expected to continue as we scale our silicon devices to the level of atomic scaling, with oxides a few atoms thick and channels with countable numbers of dopant atoms. The rise in the inherent systematic and random nonuniformity will have effects that are far-reaching in every aspect of design, manufacturing, test, and overall reliability.

It's obvious that no single research or advancement can cover all these issues and problems. In this special issue, we have included five articles that attack the subject from very different angles—namely, process monitoring, testing, adaptive circuits, and even architecture changes. We hope this diversity will give you a general idea of the proposed technologies in this field.

The first article, "Testing On-Die Process Variation in Nanometer VLSI," by Mehrdad Nourani and Arun Radhakrishnan, focuses on process monitors in the form of ring oscillators. The authors combine the results of different monitors on a chip to achieve a low-bandwidth, high-information-content mechanism to monitor process variation effects. This technique might be unconventional, but hopefully it will provoke some debates and discussions and lead to even better ideas in the future.

The second article, "Statistical Test Compaction Using Binary Decision Trees," by Sounil Biswas and Ronald (Shawn) Blanton, advocates using statistical techniques to cross-correlate different analog tests and find the minimum set required to guarantee product quality while reducing test time and cost. Traditionally, analog designs must deal with process variation because many analog designs are far more process dependent than their digital counterparts. Analog parametric testing is necessary to ensure the product meets its specifications. Analog product test costs often exceed the cost of silicon, so this study exemplifies one aspect of the use of statistical techniques.

The third article, "A DFT Approach for Testing Embedded Systems Using DC Sensors," by Soumendu Bhattacharya and Abhijit Chatterjee, addresses process variation from a different angle. Rather than testing the parameters themselves, the authors choose other, easi-

er-to-implement tests to correlate with the actual parameters. They call this test technique *alternate test*, and this article exemplifies this approach. To make their approach even better, they incorporate an on-chip sensor that can potentially provide the process and environmental feedback needed to mitigate the effect of process variation.

The fourth article, "Using Adaptive Circuits to Mitigate Process Variations in a Microprocessor Design," by Eric Fetzer, is a case study of the Montecito processor (a member of the Itanium processor family) and its use of adaptive circuits to combat process variation. Montecito is probably the biggest high-volume manufacturable die the industry has ever made, with over 1.7 billion transistors. Its designers definitely must figure out how to deal with various process variation effects, especially those on chip. It's not surprising then that they must resort to adaptive clock and other adaptive circuits.

Finally, "ElastiC: An Adaptive Self-Healing Architecture for Unpredictable Silicon," by Dennis Sylvester, David Blaauw, and Eric Karl, describes a world in which everything is pushed to the extreme: a multiple-core processor subjected to huge process variations and operating under an environment where transistors can degrade at varying rates and devices can fail or wear out prematurely. Such a world certainly would be a challenge for any designer, and a new architecture like the one described here would be called for. Is this the architecture of the future? Only time will tell.

**WE REALIZE** it's impossible in a single issue to discuss all the problems caused by process variation or the entire spectrum of solutions from all the domains. So, we've attempted to present a representative sample of the exciting ideas in this area. Hopefully, this special issue will inspire you to help solve this growing problem, which the entire industry is facing. If you think you can do better than some of the ideas presented here, we

have achieved our goals. Better yet, we'd love to be able to see new ideas and new research reported back to *D&T* as a result of this special issue. ■



**T.M. Mak** is a senior research scientist at Intel. His research interests include defect-based testing, fault effects of nanometer technology, circuit-level and physical-design test issues, IO interface and analog testing, and fault-tolerant and online testing. Mak has a BS in electrical engineering from Hong Kong Polytechnic University. He is a senior member of the IEEE.



**Sani Nassif** manages the Tools and Technology Department at IBM Austin Research Laboratory. His research interests include design and technology coupling, especially model-to-hardware matching, simulation and modeling, physical design, statistical modeling, and statistical technology characterization. Nassif has a BS in electrical engineering from the American University of Beirut and a PhD in electrical and computer engineering from Carnegie Mellon University. He is a senior member of the IEEE.

■ Direct questions and comments about this special issue to T.M. Mak, 2200 Mission College Blvd., SC12-604, Santa Clara, CA 95051, [t.m.mak@intel.com](mailto:t.m.mak@intel.com); or Sani Nassif, Tools and Technology Dept, IBM Austin Research Laboratory, 11501 Burnet Rd., MS 904-6G021, Austin, TX 78758, [nassif@us.ibm.com](mailto:nassif@us.ibm.com).

**For further information on this or any other computing topic, visit our Digital Library at <http://www.computer.org/publications/dlib>.**