



# Variation-tolerant design

**PRADIP BOSE**

[pbose@us.ibm.com](mailto:pbose@us.ibm.com)

.....As the guest editors and I introduce this year's Hot Chips theme issue, the frequency slowdown trend that is upon the industry as a result of the CMOS technology outlook has to be the single major point that stands out. It is not just the per-chip power dissipation envelope that forces this trend, although that factor alone is perhaps the major deterrent to frequency escalation at prior (historical) rates. The basic uncertainty and variance of technology-related device and interconnect parameters (both across and within individually fabricated chips) are also causing a major rethinking of the design stack—from high-level microarchitecture to physical design and layout.

There are major implications in terms of design tools support for future, technology-aware microprocessor development. Designing chips to meet worst-case timing constraints in this new world of variation-prone fabrication would in itself lead to a lower rate of frequency growth. This, compounded with power-related problems, would tend to make further technology scaling irrelevant from a raw performance viewpoint. Hence, tools and methods to aid variation-tolerant design is a new area of R&D that deserves special attention.

One key aspect of variation-tolerant design is the new attention paid to statistical timing tools. In deterministic timing analysis, you must add on guard bands to timed critical paths in a given logic stage to account for variability, thereby making it harder to meet aggressive cycle time constraints in next-generation designs. In statistical timing, designers can use nominal delays of the critical paths to fix the cycle time estimate, thereby easing the difficulty of meeting next-generation target frequencies. Manufacturers can then use post-silicon sorting and binning to yield a range of chips that operate across the various frequencies suitable for diverse markets.

Another possibility is to design for a cycle time target that the vast majority of computing scenarios would meet. However, the variation-tolerant design would have built-in checks and redundant (slow) computing paths to detect and recover from timing-related functional errors that manifest in rare corner case situations. Dynamic adaptation of processor voltage and/or frequency settings in response to workload and temperature variations, or simply fabrication differences, is yet another approach. The problem of how to make such variation- and error-tolerant designs area-efficient, while enabling steady per-

formance growth in future technology generations, is currently an active area of research at various universities. The design community wants to benefit from that research as soon as possible.

This year, Bill Dally and Keith Diefendorff have guest edited the Hot Chips special issue. Their guest editors' introduction covers the theme articles selected for publication in *IEEE Micro* in some detail. Coping with problems of variability and/or power dissipation is evident in several of the articles. The trend towards multicore, lower-frequency designs as a response to the power problem is apparent. In addition, chips such as Montecito address some of the variability concerns. Specialized accelerators to speed up key components of an application workload in a power-efficient manner are also an increasing trend, as reflected in some of these articles.

I hope this issue fulfills our aim of bringing you up to date with some of the latest advances in microprocessor designs. Enjoy!

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