RELIABILITY-AWARE MICROARCHITECTURE DESIGN

Vijay Janapa Reddi
University of Texas at Austin

How do we design error-tolerant processors and associated systems that meet historically established high reliability standards, without exceeding power budgets and cost constraints in the commodity market? This is the technological challenge facing present-day and future processor architects. In the late CMOS era, device-scaling trends have increased awareness of the various sources of unreliability at the component level. Building robust processors has become increasingly challenging in the face of growing device susceptibility to transient and hard errors. Some solutions, such as those that circumvent the power problem today, have in fact been shown to worsen conditions for the emerging new device "reliability wall." As such, future systems require processor architects to understand and integrate adaptive design techniques, at both the device and system levels, in order to ensure robust and resilient execution. The architecture layer, which includes both hardware and software, must learn to dynamically identify, detect, and recover from errors in the field. To this end, this *IEEE Micro* special issue features seven articles that focus on a set of three main topics that must be explored simultaneously to make reliability-aware design practical: understanding the tradeoffs between power, performance, and reliability; cross-layer research; and frameworks for conducting reliability tradeoff analysis.

A good number of articles in this issue focus on the interplay between power, performance, and reliability. For instance, in "Coping with Parametric Variation at Near-Threshold Voltages," Ulya R. Karpuzcu et al. discuss the impact of parameter variations on near-threshold-voltage-based computing (NTC), which has long been touted as a promising strategy to improve energy efficiency by reducing the supply voltage while exploiting parallelism to compensate for frequency loss. However, the authors demonstrate that NTC's potential cannot be harnessed without first addressing the impact of parameter variations on low-voltage operation.

Similarly, in "Improving Throughput of Power-Constrained Many-Core Processors Based on Unreliable Devices," Hao Wang and Nam Sung Kim demonstrate that by scaling back on device-level redundancy for reliability in a many-core processor, it is possible to achieve better energy efficiency, because a many-core processor can inherently tolerate defective cores by design. Keeping an eye to the future, the authors demonstrate the tradeoff between reliability and energy efficiency in the context of carbon nanotube-based devices.

The tradeoffs between redundancy and energy efficiency, however, can scale all the way from the core level down to the individual data path bit slices. For instance, in "Resilient High-Performance Processors with Spare RIBs," David J. Palframan et al. show that by adding a single bit slice and using bit-sliced redundancy along the data path, we can tolerate data path defects and improve performance in the presence of within-die

Published by the IEEE Computer Society

0272-1732/13/$31.00 © 2013 IEEE

random variations. Additionally, we can boost performance by avoiding excessively slow critical paths that are created because of random variations.

The importance of optimizing for the interplay between power, performance, and reliability is perhaps best captured by the article "Active Guardband Management in Power7+ to Save Energy and Maintain Reliability" by Charles R. Lefurgy et al. In this article, the authors demonstrate the benefits of using feedback controllers to reduce the guardband that guarantees proper circuit timing during worst-case voltage droop events in order to achieve better energy efficiency. Using real Power7 measurements, the authors demonstrate that power consumption can be automatically reduced by 11 percent for SPEC CPU2006 workloads with a negligible loss in performance and no loss in reliability.

An important and emerging tenet for reliability-aware design is that manufacturing technology alone can no longer be responsible for guaranteeing device reliability. Instead, all layers of the system stack must play a conscious role to mitigate excessive reliability overheads. Therefore, in "A Cross-Layer Technology-Based Study of How Memory Errors Impact System Resilience," Veit B. Kleeberger et al. discuss the implications of distributing resiliency tasks across all levels of the design hierarchy. Specifically, they study single-event upsets in the data cache of an autonomous robot along with the impact of voltage variations to demonstrate that resiliency solutions exist at all levels of the hardware and software stack.

In order to support reliability-aware design at all levels of the system stack, we require frameworks that let us conduct tradeoff analysis. As a step in the right direction, in "Evaluating Overheads of Multibit Soft-Error Protection in the Processor Core," Lukasz G. Szafaryn et al. present a framework for architecture-level tradeoff analysis for various soft-error protection overheads at a component and circuit level in terms of processor area, delay, and energy.

Beyond architecture design, it is important to develop frameworks that test the processor's robustness under a variety of operating conditions. In "Automating Stressmark Generation for Testing Processor Voltage Fluctuations," Youngtaek Kim et al. focus on automatically testing a processor's resilience to rapid current changes. Sudden current fluctuations can cause timing errors in the microprocessor.

Many of the articles in this issue outline future directions in reliability-aware microprocessor design. In the future, processor design will require significant collaborative effort between circuits, architecture, and firmware design to guarantee robustness. Although these articles highlight the challenges ahead of us, we look forward to seeing progress in all areas over the next several years. MICRO

Acknowledgments

Reliability-aware design requires us to rethink how we build future commodity microprocessors, and as such I believe this *IEEE Micro* special issue is timely. I thank Erik Altman for inviting me to be the guest editor for this special issue. I am grateful for his feedback and patience during the entire process, which involved soliciting articles, inviting reviewers, and gathering their feedback. Dr. Altman and I thank the authors that submitted to this special issue and all the reviewers for providing thoughtful and useful feedback on the articles.

Vijay Janapa Reddi is an assistant professor in the Department of Electrical and Computer Engineering at the University of Texas at Austin. His research interests focus on the interactions between hardware and software to explore new opportunities and synergies for cross-layer solutions that improve processor- and system-level power, performance, and reliability. Dr. Janapa Reddi has a PhD in computer science from Harvard University.

Please direct questions and comments about the articles in this special issue to Vijay Janapa Reddi, 2501 Speedway, ENS 530, University of Texas at Austin, C0803, Austin, TX 78712; vj@ece.utexas.edu.

cn *Selected CS articles and columns are also available for free at http://ComputingNow.computer.org.*