RETROSPECTIVE:
The Case for Lifetime Reliability-Aware Microprocessors

Jayanth Srinivasan[1], Sarita Adve[2], Pradip Bose[3], Jude A. Rivers[4]
[1]Stitchflow, Inc. [2]University of Illinois at Urbana-Champaign [3]IBM Research [4]Independent Consultant

I. Introduction

The processor architecture community had been well aware of the so-called "power wall" challenge since the very end of the 20th century. As such, research papers centered around the power-aware (and later the temperature-aware) microarchitecture themes started to appear in major architecture forums since 2000 (e.g., refer to a few of the earliest articles [1-5]).

At IBM Research, where the legendary Robert Dennard was still an active researcher, the architecture group worked with technology experts to understand the implications of Dennard's scaling law as we entered the 21st century. The conclusions drawn in terms of escalating power densities (and temperature) [2] led us immediately to worry about lifetime reliability (e.g., in terms of failure mechanisms like electromigration (EM), stress migration (SM) and thermal cycling (TC). The fact that the supply voltage (V_{dd}) could not scale down at historical rates any more led us further to issues beyond current densities and chip temperature: e.g., oxide breakdown (TDDB: time-dependent dielectric breakdown).

As in the preceding pioneering work on power- and temperature-aware microarchitecture [1-5], our research in the domain of reliability-aware microarchitecture started with an attempt to model the effects of the above-mentioned failure mechanisms at the processor architecture level. The work started in 2003, when the first author (Jayanth Srinivasan) began a summer internship at IBM T.J. Watson Research Center. The progress, in close consultation with reliability physics researchers at IBM (e.g., James Stathis, Sufi Zafar, C. K. Hu, E. Y. Wu, Ann Swift and several others) was rapid and the early results were stunning in that the modeled processor lifetime reliability degradation during the impending "late CMOS" design era threatened to be a veritable show-stopper! The summer project extended into a longer-term coop assignment, so that we could pursue the research to a definitive checkpoint, where the findings could be better understood, calibrated and then published. This collaborative research between industry and academia resulted in a two-part publication: (a) one that provided a general mode-driven outlook about the effect of non-ideal technology scaling on processor lifetime reliability [6]; and (b) the subject paper [7] that additionally also introduced the notion of dynamic reliability management (DRM) – a technique where the processor can respond to changing application behavior to maintain its lifetime reliability target.

II. Summary of Contributions

In retrospect, the contributions of the subject paper [7] went well beyond the ones claimed in the paper. Let us begin with the main claims that were formally listed in the paper, namely:

- Development and open-source release of an architecture-level model and a specific implementation thereof, called RAMP.
- Introduction of the notion of Dynamic Reliability Management (DRM).

Indeed, the developed model (RAMP) was quickly made available for general use in conjunction with cycle-accurate processor power-performance-temperature simulators of the day. For our work, we used IBM's homegrown PowerTimer [4], used in conjunction with IBM's Turandot performance simulator (that modeled a POWER4 processor) to provide the basic power-performance simulation capability.

The then-recent HotSpot thermal model [3] was used to project on-chip temperature values. The lifetime reliability results, derived for the SPEC benchmark suite of the day showed (see the companion paper [6]) that technology scaling from 180nm over three generations to 65nm could result in an alarming 316% increase in failure rate, on average. Of course, the modeling exercise assumed an unchanged POWER4-like design that simply got remapped successively across technology generations. Although this was admittedly an overly pessimistic view, the main point of the modeling was to show that straightforward reuse/remap of a fixed microarchitecture design was no longer going to be a scalable proposition. Significant changes in the low-level design, supported also by progressively more power-efficient micro-architectural paradigms would be needed to maintain target FIT (failures in time) rates.

Coupled with separate (prior) findings about how single processor pipeline depths (and operating clock frequencies) had reached their scaling limits [8], this class of work pointed to significant implications about future microarchitecture design. Within IBM and other processor design groups, indeed the conclusions from these papers (and others in that era) led to the realization that the days of sustained core-level frequency increases were virtually over, and instead, multi-core, multi-threaded (SMT) designs (at ~constant clock frequency) were in vogue for the foreseeable future. The DRM innovation in the subject paper [7] pointed to a solution space that architects could migrate to, in case competitive pressures forced industry to keep pushing for deeper pipelines, higher frequencies (coupled with non-ideal voltage scaling), wider

super scalar engines – which would collectively lead to higher temperatures, current densities and across-dielectric electric fields than what could be tolerated in a nominally constant-reliability regime.

III. Follow-on Research and Development Impact

The RAMP-driven predictive analysis and the DRM proposal spawned off a number of different research endeavors across academia and industry. The taxonomy in the subject DRM paper [7] already talked about soft and hard errors – both of which were of rising concern in the late CMOS, post-Dennard-scaling design era. The lifetime reliability models became more comprehensive, where aging-related effects (e.g., NBTI/PBTI) were brought into consideration. The effects of aggressive power management, coupled with late-CMOS era technology scaling with associated aging and other failure modes have raised the specter of silent data corruption (SDC) and so-called "mercurial cores" that have been reported recently by the hyper scalers (e.g., Google and Facebook). As such, RAMP-like modeling and associated DRM-style mitigation solutions may indeed need to be pursued in future work.

For the senior authors on the subject paper, the groundbreaking work represented by this industry-academic collaboration served as the impetus for a long and fruitful collaborative relationship that continues today, as well as a major focus on reliability throughout their ensuing research careers. This work has included models, metrics, open source tools, and mitigation techniques for a variety of reliability concerns, using similarly cross-stack approaches ranging from circuit level to architecture to application software (e.g., [9-22]).

We are grateful and humbled to receive this recognition for our early work in what became an important area of research in our community.

References

[1] D. Brooks, V. Tiwari, M. Martonosi, "Wattch: a framework for architectural-level power analysis and optimizations," Proc. 27th Ann. Int'l. Symp. on Comp. Arch. (ISCA), June 2000.

[2]. D. Brooks, P. Bose, S. Schuster, H. Jacobson, P. Kudva, A. Buyuktosunoglu, J-D. Wellman, V. Zyuban, M. Gupta, P. Cook, "Power-Aware Microarchitecture: Design and Modeling Challenges for Next-Generation Microprocessors," IEEE Micro 20(6): 26-44 (2000)

[3] Kevin Skadron, Mircea R. Stan, Wei Huang, Sivakumar Velusamy, Karthik Sankaranarayanan, David Tarjan, "Temperature-aware microarchitecture," Proc. 30th Ann. Int'l. Symp. on Computer Architecture (ISCA), June 2003, pp. 2-13.

[4] D. Brooks, P. Bose, V. Srinivasan, M. Gschwind, P. Emma, M. Rosenfield, "New methodology for early-stage, microarchitecture-level power-performance analysis of microprocessors," IBM J. Res. Dev. 47(5-6): 653-670 (2003).

[5] J. Srinivasan, S. V. Adve, "Predictive Dynamic Thermal Management for Multimedia Applications," Proc. 17th Annual ACM International Conference on Supercomputing (ICS '03) June 2003, 109-120.

[6] J. Srinivasan, S. Adve, P. Bose, J. Rivers, "The impact of scaling on processor lifetime reliability," Proc. Int'l. Conf. on Dependable Systems and Networks (DSN), June 2004.

[7] J. Srinivasan, S. Adve, P. Bose, J. Rivers, "The case for lifetime reliability-aware microprocessors," Proc. 31st Int'l. Symp. on Computer Architecture, June 2004.

[8] V. Srinivasan, D. Brooks, M. Gschwind, P. Bose, V. Zyuban, P. Strenski, P. Emma, "Optimizing pipelines for power and performance," *Proc. Int'l. Symp. on Microarchitecture (MICRO), 2002*, pp. 333-344.

[9] X. Li, S. V. Adve, P. Bose, and J. A. Rivers, "SoftArch: An Architecture Level Tool for Modeling and Analyzing Soft Errors," *Proc. Intl. Conf. on Dependable Systems and Networks (DSN),* June 2005, 496-505.

[10] M.-L. Li, P. Ramachandran, S. K. Sahoo, S. V. Adve, V. S. Adve, and Y. Zhou, "Understanding the Propagation of Hard Errors to Software and Implications for Resilient System Design," *Proc. Intl. on Architectural Support for Programming Languages and Operating Systems (ASPLOS),* March 2008.

[11] P. Ramachandran, S. V. Adve, P. Bose, J. A. Rivers, "Metrics for Architecture-Level Lifetime Reliability Analysis*," Proc. IEEE Intl. Symp. on Performance Analysis of Systems and Software (ISPASS),* April 2008, 202-212.

[12] X. Li, S. V. Adve, P. Bose, J. A. Rivers, "Online Estimation of Architectural Vulnerability Factor for Soft Errors," *Proc. 35th Intl. Symp. on Computer Architecture (ISCA),* June 2008.

[13] M.-L. Li, P. Ramachandran, U. R. Karpuczu, S. K. S. Hari, and S. V. Adve, "Accurate Microarchitecture-Level Fault Modeling for Studying Hardware Faults," *Proc. 15th Intl. Symp. on High-Performance Computer Architecture (HPCA),* 2009.

[14] S. K. Hari, M.-H. Li, P. Ramachandran, B. Choi, and S. V. Adve, "mSWAT: Low-Cost Hardware Fault Detection and Diagnosis for Multicore Systems," *Proc. 42nd Symp. on Microarchitecture (MICRO),* December 2009.

[15] R. Venkatagiri, A. Mahmoud, S. K. S. Hari, and S. V. Adve, "Approxilyzer: Towards A Systematic Framework for Instruction-Level Approximate Computing and its Applications to Hardware Resiliency," Proc. *49th Intl. Symp. on Microarchitecture (MICRO),* October 2016.

[16] A. Mahmoud, S. Hari, C. W. Fletcher, S. V. Adve, C. Sakr, N. Shanbhag, P. Molchanov, M. B. Sullivan, T. Tsai, and S. W. Keckler, "Optimizing Selective Protection for CNN Resilience," *Proc. 32nd Intl. Symp. on Software Reliability Engineering (ISSRE),* October 2021.

[17] J. A. Rivers, P. Bose, P. Kudva, J-D. Wellman,, P. Sanda, E. Cannon, L. Alves, "Phaser: Phased methodology for modeling the system-level effects of soft errors," *IBM J. Res. Dev.* 52(3): 293-306 (2008).

[18] J. Shin, V. Zyuban, P. Bose, T. Pinkston, "A Proactive Wearout Recovery Approach for Exploiting Microarchitectural Redundancy to Extend Cache SRAM Lifetime," *Proc. 35th Int'l. Symp. on Computer Architecture (ISCA),* June 2008, pp. 353-362.

[19] K. Swaminathan, N. Chandramoorthy, C-Y. Cher, R. Bertran, A. Buyuktosunoglu, P. Bose, "BRAVO: Balanced Reliability-Aware Voltage Optimization," *Proc. Int'l. Symp. on High Performance Computer Architecture (HPCA),* Feb. 2017, pp. 97-108.

[20] S. Mitra, P. Bose, E. Cheng, C-Y. Cher, H. Cho, R. Joshi, Y. M. Kim, C. Lefurgy, Y. Li, K. Rodbell, K. Skadron, J. Stathis, L. Szafaryn, "The resilience wall: cross-layer solution strategies," *Proc. VLSI-DAT,* 2014, pp. 1-11.

[21] N. Chandramoorthy, K. Swaminathan, M. Cochet, A. Paidimarri, S. Eldridge, R. Joshi, M. Ziegler, A. Buyuktosunoglu, P. Bose, "Resilient Low Voltage Accelerators for High Energy Efficiency*," Proc. Int'l. Symp. on High Performance Computer Architecture (HPCA),* 2019, pp. 147-158.

[22] J. Leng, A. Buyuktosunoglu, R. Bertran, P. Bose, Q. Chen, M. Guo, V. J. Reddi, "Asymmetric Resilience: Exploiting Task-Level Idempotency for Transient Error Recovery in Accelerator-Based Systems," *Proc. Int'l. Symp. on Computer Architecture (HPCA),* 2020, pp. 44-57.