A Comprehensive Timing, Power, Thermal, and Reliability Model for Exascale Node Architectures

In modern processors, the interactions between performance, power, thermal, and reliability are complex. Neglecting one aspect, such as reliability, might lead to unrealistically optimistic evaluations of the other aspects due to the (implicit) assumption on non-faulty devices [15]. The sheer size and the compute, power, and mean-time-to-intervention (MTTI) requirements of an exascale system further increase the need to model the dynamic interactions between all four aspects [5]. Hence a detailed analysis to determine the capability of a processing node must precede extrapolations to the datacenter using other methodologies [6].

Along with accuracy, detailed cycle-level simulations provide flexibility to simulate future programming models and associated hardware support that will exist in the exascale timeframe. As an example, Heterogeneous System Architecture (HSA) [8] promises full coherence between the CPU and GPU for more efficient GPU compute off-load in future accelerated processing unit (APU) systems. Only by modeling functionality and timing, can designers and programmers understand HSA's impact.

To model the detailed interactions while providing a flexible infrastructure, we propose a comprehensive simulation flow using a combination of three existing open-source simulators: gem5 [2], McPAT [5], and HotSpot [12], as well as a gem5 compatible Architecture Vulnerability Factor (AVF) model for modeling reliability [13]. These simulators have an impressive history of success with the simulators being cited in 131, 297, and 269 publications respectively. As evidence by their citation count, these models are mature, but to our knowledge, no one has tightly integrated them into a single closed-loop methodology.

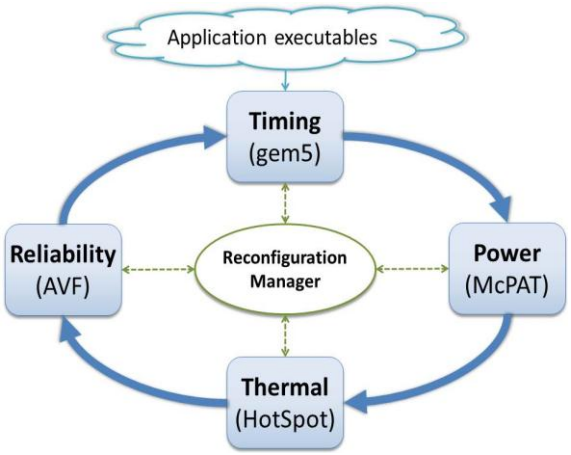

Figure 1: Proposed simulation methodology

Figure 1 illustrates our general approach where data is passed between models and a reconfiguration manager uses the models' outcomes to dynamically adjust the system configuration. The methodology begins by performing cycle-level execution-driven simulation of an application on gem5, including any advanced programming model features such as HSA. Periodically the gem5 timing simulation is interrupted and its output data feeds into McPAT to estimate the node power. The McPAT calculated node power then feeds into Hotspot to determine the node's thermal characteristics. Based on the outputs of gem5, McPAT, and Hotspot, the AVF component calculates the soft error vulnerability. This reliability analysis can be used dynamically stop the simulation if too many errors are detected or post-processing can determine what RAS mechanisms are best suited to meet reliability targets. Finally, before starting the next time step, the reconfiguration controller adjusts voltage, frequency or other parameters to maintain operation within set constraints.

The proposed methodology provides the necessary mechanisms to model the fine-grain

and multidimensional interactions of tightly integrated node components. For example the evolution of interposer technology [4] suggests that exascale nodes will likely not only integrate heterogeneous processing resources (such as a CPU and GPU) into the same package, but they will also include integrated network interfaces, as well as stacks of DRAM. In such dense designs the interaction between components will not be limited to data communication. Instead, the dynamic sharing of a fixed power envelope between components will also have first-order impact on performance. Furthermore, the thermal qualities of the design impact leakage power, as well as the reliability and wear-out of logic [14] and memory (both DRAM and NVRAM [16]). Similarly, the voltage supplied to particular components impacts their reliability and thus certain components can be turned off in order to raise the voltage and increase the reliability of other components [3]. The proposed methodology allows us to model all of these important interactions and tradeoffs.

Accurate projections in to the exascale era, precise management of interactions within the node can only be achieved if the individual simulation components are accurate themselves. Hence we propose a closed loop calibration for individual components based on data from real hardware that will minimize the error caused by improving microarchitecture and technology. This calibration may occur at different levels. A corporate researcher may choose to use internal data, while academics may use techniques such as micro-benchmarking on real hardware [5].

Assessment Criteria

Challenges addressed: The proposed methodology models node-level interactions of power, reliability, thermal, and performance with a dynamic reconfiguration manager.

Maturity: These individual models have already been successfully used in various evaluations. By integrating the models into one comprehensive model, we can substantially improve their usability and overall impact.

Uniqueness: Our proposed comprehensive model helps evaluate the challenging performance, power, and reliability requirements that exascale systems are facing first before other smaller-scale systems. In addition, our model infrastructure includes an APU, rather than just a GPU or a CPU, enabling explorations of future programming models like HSA.

Novelty: To our knowledge, no one has integrated all four models into one comprehensive infrastructure or included a dynamic reconfiguration manager. We believe this tight integration will improve the overall accuracy.

Applicability: The prior successes of the individual models suggest that the broad community (from mobile to servers) would benefit tremendously from the combined methodology. Furthermore, using these open source simulators, licensed under industry friendly licenses, simplifies the distribution and usage of our combined methodology.

Effort: While the individual components already exist, methodologies this complex require continual maintenance and correlation. A general power methodology is critical to allow academic researchers to create coarse-grain empirical models while industrial researchers use the framework to build fine-grain empirical models from proprietary data. Furthermore, increasing the length of simulated time is important to capture power, thermal, and reliability, though it is challenging with current single-threaded interpretation-based simulation. One solution is to parallelize our simulation engine using multiple event queues. Another solution is to dynamically switch or sample between a fast functional model, such as QEMU [11], KVM [9], or HP COTSon [1], and the detailed, cycle-level gem5 simulation.

References

1. Argollo, Eduardo; Falcón Ayose; Faraboschi, Paolo; Monchiero, Matteo; Ortega, Daniel. "COTSon: Infrastructure for Full System Simulation", *ACM SIGOPS Operating Systems Review*, January 2009.
2. Binkert, Nathan and Beckmann, Bradford and Black, Gabriel and Reinhardt, Steven K. and Saidi, Ali and Basu, Arkaprava and Hestness, Joel and Hower, Derek R. and Krishna, Tushar and Sardashti, Somayeh and Sen, Rathijit and Sewell, Korey and Shoaib, Muhammad and Vaish, Nilay and Hill, Mark D. and Wood, David A. "The gem5 simulator", *SIGARCH Computer Architecture News*, May 2011.
3. Calimera, Andrea; Macii, Enrico; Poncino, Massimo. "NBTI-Aware Power Gating for Concurrent Leakage and Aging Optimization", In Proc. of the *14th ACM/IEEE International Symposium on Low Power Electronics and Design (ISLPED)*, 2009.
4. Clermidy, Fabien; Darve, Florian; Dutoit, Denis; Lafi, Walid; Vivet, Pascal. "3D Embedded Multi-Core: Some Perspectives", In Proc. *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, March 2011.
5. J. Daly et al. Inter-Agency Workshop on HPC Resilience at Extreme Scale. National Security Agency, Advanced Computing Systems, February 21-24, 2012.
6. Hendry, Gilbert and Rodrigues, Arun. SST: A simulator for exascale co-design. In *ASCR/ASC Exascale Research Conference*, April 2012.
7. Hong, Sunpyo and Kim, Hyesoon. "An integrated GPU power and performance model", In Proceedings of the *37th annual international symposium on Computer architecture (ISCA)*, June 2010.
8. HSA Foundation, AFDS 2011 Phil Rogers Keynote. "The Programmer's Guide to the APU Galaxy", June 2012.
9. Kivity, Avi; Kamay, Yaniv; Laor, Dor; Lublin, Uri; Liguori, Anthony. "kvm: the Linux Virtual Machine Monitor", In Proc. of the *Linux Symposium*, July 2007.
10. Li, S.; Ahn, J. H.; Strong, R. D.; Brockman, J. B.; Tullsen, D. M.; Jouppi, N. P. "McPAT: An integrated power, area, and timing modeling framework for multicore and manycore architectures", In Proc. the *42nd Ann. IEEE/ACM International Symp Microarchitecture (MICRO)*, Dec. 2009.
11. Shen, Shye-Tzeng; Lee, Shin-Ying; Chen, Chung-Ho. "Full System Simulation with QEMU: An Approach to Multi-View 3D GPU Design", In Proc. of *IEEE International Symposium on Circuits and Systems (ISCAS)*, May 2010.
12. Skadron , K.; Stan, M. R.; Huan, W.; Velusamy, S.; Sankaranarayanan, K.; Tarjan, D. "Temperature-Aware Microarchitecture." In *Proceedings of the 30th International Symposium on Computer Architecture*, June 2003.
13. Sridharan, Vilas and Kaeli, D. R. "Using Hardware Vulnerability Factors to Enhance AVF Analysis". In Proceedings of the *37th annual international symposium on Computer architecture (ISCA)*, June 2010.
14. Srinivasan, Jayanth; Adve, Sarita V.; Bose, Pradip; Rivers, Jude A. "Lifetime Reliability: Toward an Architectural Solution", *IEEE Micro*, May-June 2005.
15. Stevens, R. et al., "Scientific Grand Challenges: Architectures and Technologies for Extreme Scale Computing", *Workshop Report*, December 2009.
16. Yoon, Doe Hyun; Erez, Mattan. "Virtualized ECC: Flexible Reliability in Main Memory", *IEEE Micro*, Jan-Feb 2011