Models and Techniques for Green High-Performance Computing

Vignesh Adhinarayanan
Department of Computer Science, Virginia Tech
Blacksburg, Virginia 24061
avignesh@vt.edu

ABSTRACT

Modern high-performance computing (HPC) systems are power-limited. For instance, the U.S. Department of Energy has set a power envelope of 20MW for the exascale supercomputer expected to arrive in 2022-23. Achieving this target requires a 10.8-fold increase in performance over todayfis fastest supercomputer with only a 1.3-fold increase in power consumption. As a consequence, the architecture of an HPC system is changing rapidly (e.g., via heterogeneity and hardware overprovisioning). In my dissertation, I will address (i) modeling, (ii) management, and (iii) evaluation challenges concerning power and energy in this changing landscape of HPC systems.

Accurate power measurement is essential for the efficient functioning of a power management system. Years of research has yielded methodologies for developing accurate, high-resolution power proxies for CPUs. However, such methodologies rely on the availability of numerous hardware performance counters which are limited in number in newer accelerators (e.g., GPUs). To address this limitation, we use data from a low-resolution power meter to refine performance-counter-based power proxies at runtime to achieve both high accuracy and high resolution for GPU power measurement [7]. While adopting the best practices from CPU power modeling results in an average error of 6%, our online power modeling approach reduces the error to nearly 1%.

On-chip data movement is considered to be a major source of power consumption in modern processors. Properly understanding the power that applications consume in moving data is vital for inventing mitigation strategies. Past research on measuring data movement power on real hardware failed to distinguish data movement power from data access power. In this work, I developed a novel approach based on the physical distance of data movement to measure interconnect power accurately and study its characteristics [6]. Our evaluation shows that up to 14% of the dynamic power is consumed by the interconnect (which is less than what previous studies have suggested). However, our projection also shows that this fraction is expected to increase in the future. Therefore, we also present a range of mitigation strategies that can reduce the interconnect power in the future.

To evaluate any idea in the systems area, we need to choose test cases carefully. To systematically select a concise, but useful, set of applications for evaluating the techniques, I employ statistical methods such as principal component analysis (PCA) and hierarchical clustering [1]. Our evaluation shows that all four benchmark suites studied (namely, Parboil, SHOC, Rodinia, and SPEC ACCEL) contains significant redundancy, and we can perform a thorough evaluation using only a fraction of the applications in these benchmark suites.

A major goal of this work is to maximize an application's performance under some power (or energy) constraint. Our broad solution here is to (i) identify the architectural component acting as the power (or energy) bottleneck and (ii) alleviate the bottleneck by making more power (or energy) available to these components. The re-allocation of power budgets to the various components can occur either statically or dynamically. Static reallocation has already been explored [3–5]. The problem of dynamically reallocating power budget is currently being studied as explained next.

Several runtime solutions for power management require "reconfiguration" that can occur only at certain time intervals (e.g., some implementations of DVFS). Such restrictive power management techniques require proactive planning for the future. To do so, I devised three phase predictors that look for exact and approximate patterns and trends during an applicationfis execution and forecast good configurations for the future phases. I also developed a meta-predictor to choose the best of the three predictors for a given application. Using these forecasting schemes, I am working on developing a power management solution that improves the performance of an application under a power budget [2, 8].

ACM Reference format:
Vignesh Adhinarayanan. 2018. Models and Techniques for Green High-Performance Computing. In *Proceedings of ICPP '18: International Conference on Parallel Processing, Eugene, Oregon, USA, August 13–16, 2018 (ICPP '18),* 2 pages.
DOI: 10.1145/nnnnnnn.nnnnnnn

ICPP '18, Eugene, Oregon, USA

978-x-xxxx-xxxx-x/YY/MM…$15.00
DOI: 10.1145/nnnnnnn.nnnnnnn

REFERENCES

[1] Adhinarayanan, V., and Feng, W. An automated framework for characterizing and subsetting GPGPU workloads. In *2016 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)* (2016), pp. 307–317.
[2] Adhinarayanan, V., and Feng, W. CPU-GPU Power Sloshing via Hybrid Phase Prediction Techniques (In Preparation). In *ACM SIGMETRICS* (2019).
[3] Adhinarayanan, V., Feng, W., Rogers, D., Ahrens, J., and Pakin, S. Characterizing and Modeling Power and Energy for Extreme-Scale In-Situ Visualization. In *2017 IEEE International Parallel and Distributed Processing Symposium (IPDPS)* (2017), pp. 978–987.
[4] Adhinarayanan, V., Feng, W., Woodring, J., Rogers, D., and Ahrens, J. On the Greenness of In-Situ and Post-Processing Visualization Pipelines. In *2015 IEEE International Parallel and Distributed Processing Symposium Workshop (IPDPSW)* (2015), pp. 880–887.
[5] Adhinarayanan, V., Pakin, S., Rogers, D., Feng, W., and Ahrens, J. Performance, Power, and Energy of In-Situ and Post-Processing Visualization: A Case Study in Climate Simulation. In *2015 ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis (SC)* (2015). Best Research Poster Finalist.

[6] Adhinarayanan, V., Paul, I., Greathouse, J. L., Huang, W., Pattnaik, A., and Feng, W. Measuring and modeling on-chip interconnect power on real hardware. In *2016 IEEE International Symposium on Workload Characterization (IISWC)* (2016), pp. 1–11. Best Paper Award.

[7] Adhinarayanan, V., Subramaniam, B., and Feng, W. Online Power Estimation of Graphics Processing Units. In *2016 16th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)* (2016), pp. 245–254.

[8] Vignesh Adhinarayanan, Wu Feng, J. G. I. P. Proactive Power Management Policies for On-Chip Interconnect (In preparation). In *ACM SIGMETRICS* (2019).