

# Memory and Thread Synchronization Contention-Aware DVFS for HPC systems

Karunakar R. Basireddy, Eduardo W. Wachter, Bashir M. Al-Hashimi and Geoff V. Merrett  
University of Southampton  
Southampton, United Kingdom  
{krb1g15, eww1n17, bmah, gvm}@ecs.soton.ac.uk

Due to the operating costs and failure rates of computing platforms, energy efficiency has become a major concern for modern and future many-core systems. In the quest for high performance, the power consumption growth rate must slow-down while delivering more performance per unit of power. To improve the energy efficiency of such systems, processors are equipped with low-power techniques such as dynamic voltage and frequency scaling (DVFS) and power capping. These techniques must be controlled carefully as per the workload; otherwise, it may result in significant performance loss and/or power consumption due to system overheads (e.g. DVFS transition latency). Existing approaches [1], [2] are not effective in adapting to workload variations as they do not consider the combined effect of application compute-/memory-intensity, thread synchronization contention, and non-uniform memory accesses (NUMAs) owing to the underlying processor architecture.

This poster discusses a workload-aware runtime energy management technique that takes the aforementioned factors into account for efficient  $V$ - $f$  control [3]. Fig. 1 illustrates various steps in the presented approach as a flow-chart. The processor workload is measured using Memory Accesses Per Micro-operation (MAPM) and utilization, taking thread synchronization contention into account. Moreover, latency due to NUMAs is computed by monitoring the remote and local memory accesses during the application execution. To accomplish this, four hardware performance monitoring counters (PMCs) are used to calculate MAPM, utilization and NUMA latency. To determine the appropriate  $V$ - $f$  setting, a binning-based approach with two classification layers is employed which takes utilization and MAPM as inputs. Furthermore, our approach works on both per-core and system-wide DVFS




Fig. 1. Different steps in the approach presented in [3].




Fig. 2. Comparison of the various approaches for single, double and triple application scenarios in terms of energy consumption, executing on the Xeon E5-2630. Applications [4], [5]: bfs - Breadth First Search; pf - Particle Filter; nw - Needleman-Wunsch; km - K-Means; cg - Conjugate Gradient.

supporting platforms. Experimental validation is performed on the 12-core (24 threads) Intel Xeon E5-2630 and 61-core (244 threads) Xeon Phi 7620P many-core platforms. The former supports per-core DVFS, whereas the latter is based on system-wide DVFS. To illustrate benefits of the presented approach, various application scenarios from NAS [4] and Rodinia [5] benchmark suites are used. Fig. 2 gives energy consumption on the Xeon E5-2630 for various approaches. When evaluated on the Xeon E5-2630 and Phi, presented approach achieves energy savings of up to 81.2% and 60.9% compared to CONS, OD and PERF, respectively.

## ACKNOWLEDGEMENTS

This work was supported in part by the EPSRC under EP/L000563/1 and EP/K034448/1 (the PRIME Programme, [www.prime-project.org](http://www.prime-project.org)). Experimental data used in this paper can be found at <https://doi.org/10.5258/SOTON/D0547>.

## REFERENCES

- [1] M. S. Vaibhav Sundriyal, "Runtime power-aware energy-saving scheme for parallel applications," in *Iowa State University Computer Science Technical Reports*, 2015, p. 17.
- [2] A. Marathe *et al.*, "A run-time system for power-constrained HPC applications," in *International conference on high performance computing*. Springer, 2015, pp. 394–408.
- [3] B. K. Reddy *et al.*, "Workload-aware runtime energy management for hpc systems," in *International Workshop on Optimization of Energy Efficient HPC & Distributed Systems (OPTIM)*, 2018, p. 8.
- [4] H.-Q. Jin, M. Frumkin, and J. Yan, "The OpenMP implementation of NAS parallel benchmarks and its performance," 1999.
- [5] S. Che *et al.*, "Rodinia: A benchmark suite for heterogeneous computing," in *Workload Characterization, 2009. IISWC 2009. IEEE International Symposium on*. IEEE, 2009, pp. 44–54.