CRISPIER: A CRISP-Based DVFS Implementation in gem5 for AMD GPUs

Nishant Aggarwal, Rutwik Jain, Matthew D. Sinclair
University of Wisconsin-Madison
Madison, WI, US
{naggarwal28, rnjain, msinclair}@wisc.edu

I. Introduction

Power is a first-order constraint for GPU systems due to increasing GPU power limits and growing workload demands across datacenters and supercomputers [1]–[3]. Therefore, architectural simulators such as gem5 [4], [5] are important to study power–performance trade-offs and develop power-efficiency optimizations. However, this requires high-fidelity power and frequency scaling support in simulation. Unfortunately, gem5's GPU models [6] do not support Dynamic Voltage and Frequency Scaling (DVFS), limiting power-related research. In this work, we implement CRISP [7], a state-of-the-art analytical GPU DVFS model, across gem5's AMD GPU models. CRISP classifies execution cycles as elastic cycles that scale with Compute Unit (CU) frequency, and inelastic cycles that are bound by off-chip memory latency. CRISP uses this breakdown to accurately predict workload performance at various clock frequencies without having to re-run the workload. However, CRISP was originally designed for GPU architectures that have evolved since, and its approach does not translate cleanly to modern designs. We bridge this gap by designing CRISPIER, an algorithmic modification of CRISP that adapts its core principles to modern GPU architectures where CRISP's original assumptions no longer hold.

II. Implementation & Methodology

CRISP operates over a fixed interval of cycles called a *window*, collecting execution statistics to drive the next frequency decision. Beyond elastic and inelastic cycles, CRISP tracks a third category, overlapped cycles, where compute runs concurrent to the outstanding memory accesses. Because multiple misses may be in flight simultaneously, CRISP attributes inelastic cycles using the critical path [8], the longest chain of dependent misses that actually bounds the execution time.

CRISPIER adapts this classification to three architectural changes in modern GPUs. First, CRISP assumed a fixed issue width of two, while modern GPUs have multiple heterogeneous pipelines, each issuing independently. CRISPIER therefore classifies a cycle as busy based on concurrent activity across all execution resources. Second, CRISP tracks store-stall cycles, idle cycles caused by the store queue filling up. However, modern AMD GPUs do not have a dedicated store buffer, so CRISPIER removes this category. Third, CRISP assumed that the last level cache sat in a separate clock domain. However, in modern GPUs the L1 and L2 share a clock domain with the CUs, so CRISPIER uses the last level cache miss as the point of inelasticity.

On top of this classifier, we build a per-CU DVFS infrastructure where each CU is assigned its own clock and voltage domain. At every window boundary, a dedicated controller evaluates the CRISP statistics collected over that window and selects the P-state minimizing the chosen objective, either EDP or ED^2P. We extend gem5's native CPU DVFSHandler to GPUs with a wrapper that decides when and what to transition, and extend gem5's `MathExprPowerModel` [9] to enable GPU power estimation.

We evaluate CRISPIER using three micro-benchmarks spanning compute-heavy (MaxFlops), mixed (Square), and memory-heavy (MemBandwidth) workloads [10], [11]. In steady state, MaxFlops produces 97% elastic classification, MemBandwidth produces 80% inelastic classification, and Square produces a mixed profile with 65% elastic and 36% inelastic cycles, all consistent with expected behavior. The controller correctly differentiates P-state selection across workloads, with MaxFlops predominantly running at high frequency and MemBandwidth at low frequency, and shows consistent shifts within a single kernel as execution characteristics change across windows.

III. Conclusion & Future Work

We present CRISPIER, a CRISP-based GPU DVFS implementation in gem5, adapted to AMD CDNA architectures with per-CU clock and voltage domains, timestamp-based critical path tracking, and an occupancy-based cycle classifier that handles heterogeneous execution pipelines. Going forward, we plan to upstream this infrastructure to the gem5 community, replace the current power model with a more modern model such as AccelWattch [12], and fine-tune classification thresholds and DVFS policy through evaluation on standard benchmark suites.

Acknowledgments

This work is supported in part by the National Science Foundation grant ENS-1925485.

References

[1] R. Jain, Y. Jiang, M. D. Sinclair, and S. Venkataraman, "Minos: Systematically Classifying Performance and Power Characteristics of GPU Workloads on HPC Clusters," in *Proceedings of the 52nd ACM Conference on Measurement and Analysis of Computing Systems*, ser. SIGMETRICS, June 2026.

[2] J.-s. Huang, "GTC March 2025 Keynote," in *Proceedings of GPU Technology Conference*, ser. GTC, 2025.

[3] A. L. V. Solórzano, K. Sato, K. Yamamoto, F. Shoji, J. M. Brandt, B. Schwaller, S. P. Walton, J. Green, and D. Tiwari, "Toward Sustainable HPC: In-Production Deployment of Incentive-Based Power Efficiency Mechanism on the Fugaku Supercomputer," in *International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC, 2024, pp. 1–16.

[4] N. Binkert, B. Beckmann, G. Black, S. K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D. Hower, T. Krishna, S. Sardashti, R. Sen, K. Sewell, M. Shoaib, N. Vaish, M. D. Hill, and D. A. Wood, "The gem5 simulator," *ACM SIGARCH Computer Architecture News*, vol. 39, no. 2, pp. 1–7, 2011.

[5] J. Lowe-Power, A. M. Ahmad, A. Akram, M. Alian, R. Amslinger, M. Andreozzi, A. Armejach, N. Asmussen, S. Bharadwaj, G. Black, G. Bloom, B. R. Bruce, D. R. Carvalho, J. Castrillon, L. Chen, N. Derumigny, S. Diestelhorst, W. Elsasser, M. Fariborz, A. Farmahini-Farahani, P. Fotouhi, R. Gambord, J. Gandhi, D. Gope, T. Grass, B. Hanindhito, A. Hansson, S. Haria, A. Harris, T. Hayes, A. Herrera, M. Horsnell, S. A. R. Jafri, R. Jagtap, H. Jang, R. Jeyapaul, T. M. Jones, M. Jung, S. Kannoth, H. Khaleghzadeh, Y. Kodama, T. Krishna, T. Marinelli, C. Menard, A. Mondelli, T. Mück, O. Naji, K. Nathella, H. Nguyen, N. Nikoleris, L. E. Olson, M. Orr, B. Pham, P. Prieto, T. Reddy, A. Roelke, M. Samani, A. Sandberg, J. Setoain, B. Shingarov, M. D. Sinclair, T. Ta, R. Thakur, G. Travaglini, M. Upton, N. Vaish, I. Vougioukas, Z. Wang, N. Wehn, C. Weis, D. A. Wood, H. Yoon, and Éder F. Zulian, "The gem5 simulator: Version 20.0+," 2020.

[6] A. Gutierrez, B. M. Beckmann, A. Dutu, J. Gross, M. LeBeane, J. Kalamatianos, O. Kayiran, M. Poremba, B. Potter, S. Puthoor, M. D. Sinclair, M. Wyse, J. Yin, X. Zhang, A. Jain, and T. Rogers, "Lost in abstraction: Pitfalls of analyzing gpus at the intermediate language level," in *2018 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, 2018, pp. 608–619.

[7] R. Nath and D. Tullsen, "The crisp performance model for dynamic voltage and frequency scaling in a gpgpu," in *Proceedings of the 48th International Symposium on Microarchitecture*, ser. MICRO-48. New York, NY, USA: Association for Computing Machinery, 2015, p. 281–293. [Online]. Available: https://doi.org/10.1145/2830772.2830826

[8] R. Miftakhutdinov, E. Ebrahimi, and Y. N. Patt, "Predicting performance impact of dvfs for realistic memory systems," in *2012 45th Annual IEEE/ACM International Symposium on Microarchitecture*, 2012, pp. 155–165.

[9] B. K. Reddy, M. J. Walker, D. Balsamo, S. Diestelhorst, B. M. Al-Hashimi, and G. V. Merrett, "Empirical cpu power modelling and estimation in the gem5 simulator," in *2017 27th International Symposium on Power and Timing Modeling, Optimization and Simulation (PATMOS)*, 2017, pp. 1–8.

[10] B. R. Bruce, A. Akram, H. Nguyen, K. Roarty, M. Samani, M. Fariborz, T. Reddy, M. D. Sinclair, and J. Lowe-Power, "Enabling Reproducible and Agile Full-System Simulation," in *IEEE International Symposium on Performance Analysis of Systems and Software*, ser. ISPASS. Los Alamitos, CA, USA: IEEE Computer Society, 2021, pp. 183–193.

[11] C. Jamieson, A. Chandrashekar, I. McDougall, and M. D. Sinclair, "GAP: gem5 GPU Accuracy Profiler," in *4th gem5 Users' Workshop*. New York, NY, USA: Association for Computing Machinery, June 2022.

[12] V. Kandiah, S. Peverelle, M. Khairy, J. Pan, A. Manjunath, T. G. Rogers, T. M. Aamodt, and N. Hardavellas, "AccelWattch: A Power Modeling Framework for Modern GPUs," in *Proceedings of the 54th IEEE/ACM International Symposium on Microarchitecture*, ser. MICRO. New York, NY, USA: Association for Computing Machinery, 10 2021, p. 738–753. [Online]. Available: https://doi.org/10.1145/3466752.3480063