Synergy of Dynamic Frequency Scaling and Demotion on DRAM Power Management: Models and Optimizations (Supplementary File)

Yanchao Lu, Bingsheng He, Xueyan Tang, *Senior Member, IEEE,* and Minyi Guo, *Senior Member, IEEE*

Appendix A
Results on Full System Energy Consumption

In this section, we evaluate the impact of Hybrid on full-system energy consumption and performance. We start by performing back-of-envelop calculations, following previous studies [1], [2]. We assume that the average power consumption of memory system accounts for 40% of the total system power in the baseline policy (i.e., BASE), and calculate a fixed average power estimate (i.e., the remaining 60%) for all other components. Thus, the energy consumption of all other components (i.e., non-memory system energy consumption) is proportional to the program execution time, which is usually consistent with the real-world case [1], [2]. This ratio (40%) has been chosen as the current contribution of memory system to entire system power consumption [1], [2], [3], [4], [5], [6]. We also study the impact of varying this ratio in this evaluation. Architectural characteristics and experimental parameters are the same as those used in Section 5.2.

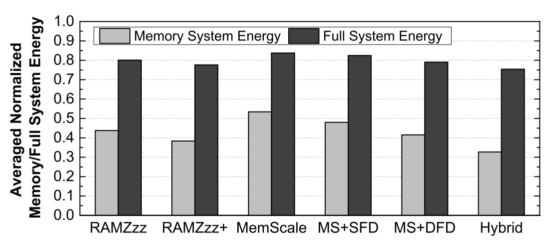

Fig. 1. Full system energy comparison.

- *Yanchao Lu and Minyi Guo are with the Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai 200240, China. The work of Yanchao Lu is done when he was a visiting student in Nanyang Technological University, Singapore. E-mail: chzblych@sjtu.edu.cn, guo-my@cs.sjtu.edu.cn.*
- *Bingsheng He and Xueyan Tang are with the School of Computer Engineering, Nanyang Technological University, Singapore 639798. Bingsheng He is the corresponding author. E-mail: bshe@ntu.edu.sg, asxytang@ntu.edu.sg.*

Figure 1 presents the averaged normalized full system energy of six approaches evaluated in this study. Hybrid is superior to BASE in both memory and full system energy savings. Hybrid's average full system energy savings is 25% on tested workloads. Compared with other power mangement polices, Hybrid still outperforms other polices on all workloads in terms of full system energy, while leads to similar performance degradations to RAMZzz+ and MemScale+DFD. Specifically, we observe that Hybrid has an average reduction of 6%, 4%, 11%, 9% and 5% over RAMZzz, RAMZzz+, MemScale, MemScale+SFD and MemScale+DFD in full system energy, respectively.

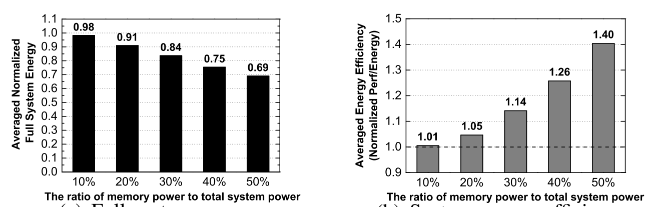

(a) Full system energy (b) System energy efficiency

Fig. 2. The impact of the ratio of memory power to total system power.

We further study the ratio of power consumption of the memory subsystem to the overall power consumption of the full system. Particularly, we vary the ratio from 10% to 50%. Figures 2(a) and 2(b) show that the fraction of memory power has a significant effect on both full system energy and system energy efficiency, respectively. We calculate the ratio of normalized performance (i.e., the reciprocal of program execution time) over energy as the system energy efficiency. A higher bar means better efficiency in Figure 2(b). Increasing the ratio from 10% to 50% (i.e., the power contribution of other components are reduced from 90% to 50%), the averaged normalized full system energy consumption of Hybrid decreases from 0.98 to 0.69, while the averaged normalized system energy efficiency increases from 1.01 to 1.40.

Appendix B
Results on ED^2 Optimizations

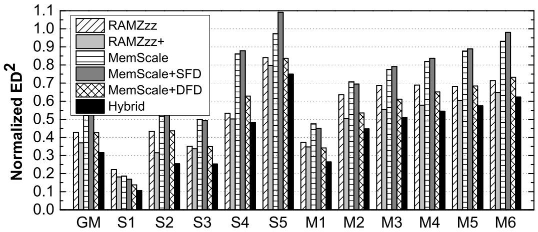

Fig. 3. Comparing ED^2 of Hybrid to other approaches with the optimization goal of ED^2.

In this section, we present some additional results for the optimization goal of ED^2. Figure 3 compares normalized ED^2 with a penalty budget of 5%. Architectural characteristics and other experimental parameters are the same as those used in Section 5.2. As can be seen, Hybrid outperforms all the other schemes. We also analyze the individual impacts of demotion (RAMZzz and RAMZzz+) and DFS (MemScale) in Hybrid. For demotion, we observe that Hybrid has average reductions of 30% and 18% in ED^2 over RAMZzz and RAMZzz+, respectively. For DFS, we observe that Hybrid has an average reduction of 45% in ED^2 over MemScale. Our proposed mechanism achieves the effectiveness and flexibility in different optimization goals.

References

[1] Q. Deng and et al., “Memscale: active low-power modes for main memory,” in *ASPLOS ’11*, 2011.
[2] N. Chatterjee, M. Shevgoor, R. Balasubramonian, A. Davis, Z. Fang, R. Illikkal, and R. Iyer, “Leveraging heterogeneity in dram main memories to accelerate critical word access,” in *MICRO ’45*, 2012.
[3] U. Hoelzle and L. A. Barroso, *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines*, 1st ed. Morgan and Claypool Publishers, 2009.
[4] J. Ousterhout and et al., “The case for ramclouds: scalable high-performance storage entirely in dram,” *SIGOPS Oper. Syst. Rev.*, vol. 43, no. 4, 2010.
[5] D. Tsirogiannis, S. Harizopoulos, and M. A. Shah, “Analyzing the energy efficiency of a database server,” in *SIGMOD ’10*, 2010.
[6] Y. Lu, D. Wu, B. He, X. Tang, J. Xu, and M. Guo, “Rank-aware dynamic migrations and adaptive demotions for dram power management,” *CoRR*, vol. abs/1409.5567, 2014.