Rank-Aware Dynamic Migrations and Adaptive Demotions for DRAM Power Management (Supplementary File)

Yanchao Lu, Donghong Wu, Bingsheng He, Xueyan Tang, Jianliang Xu and Minyi Guo

APPENDIX A
INDIVIDUAL IMPACTS

We now study the individual impact of dynamic migrations and adaptive demotions on RAMZzz with the optimization goal of ED^2 on SPEC 2006 workloads. We present the figures for DDR3 memory architecture only and comment on other architectures without figures when appropriate.

A.1 Studies on Dynamic Migrations

We study the impact of dynamic migrations, comparing RAMZzz and RZ–SP (RZ–SP uses the adaptive demotion scheme with no page migration). Figure 1 presents ED^2 results for RAMZzz and RZ–SP on DDR3.

RAMZzz has much lower ED^2 than RZ–SP, with an average reduction of 23.3%. The reduction depends on the memory footprint and memory access intensiveness. The reduction is more significant on memory-intensive workloads (such as M4) or workloads with small memory footprint (such as S3 and S4). If a workload has a small memory footprint, the page migration has a small overhead on both the delay and the energy consumption, and the portion of cold ranks is higher. If the memory access of a workload is more intensive, many idle periods are too short and RZ–SP has less opportunities for saving background power (even with our proposed adaptive demotion scheme). On those two kinds of workloads, page migration is important for the effectiveness of power management. In contrast, when the memory access is less intensive or has a large memory footprint, RZ–SP is quite competitive to RAMZzz. The example workloads include S1 and M2.

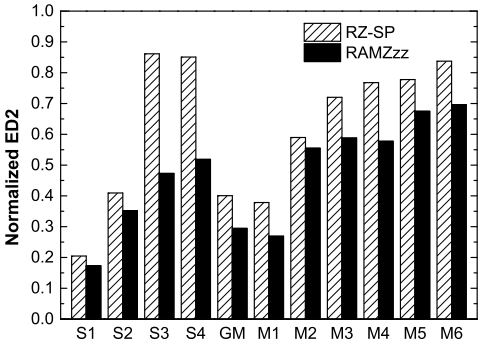

Fig. 1. Comparing ED^2 of RAMZzz and RZ–SP.

Figure 2 shows the breakdown of time stayed in different power states for RZ–SP on DDR3. Comparing with RAMZzz, the portions of time of those lower-power states with a higher resynchronization time are much smaller. RZ–SP demotes the ranks into PRE_PDN_FAST and PRE_PDN_SLOW states in most times, whereas RAMZzz demotes into even lower-power states, i.e., SR_FAST and SR_SLOW. That is because dynamic page migration is able to create longer idle periods. For example, the percentage of SR_SLOW is almost zero in RZ–SP for memory intensive workloads, such as S3, S4 and M2–4, while SR_SLOW has a significant portion in RAMZzz for those workloads. Since page migrations are disabled in RZ–SP, the total delay of RZ–SP is slightly smaller than that of RAMZzz, less than 2.5% for all workloads.

To summarize the impact of dynamic page migrations, we observe RAMZzz has much lower ED^2 than RZ–SP on three DRAM architectures, with an average reduction of 23.3%, 15.8% and 17.1%, and a range of 5.2–45.1%, 2.8–39.6% and 1.7–41.1% on DDR3, DDR2 and LPDDR2 respectively. The reduction is more significant on memory-intensive workloads or workloads with small memory footprint on DDR2 and LPDDR2.

- *Y. Lu, D. Wu and M. Guo are with the Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai 200240, China.*
- *B. He and X. Tang are with the School of Computer Engineering, Nanyang Technological University, Nanyang Ave, Singapore 639798.*
- *J. Xu is with the Department of Computer Science, Hong Kong Baptist University, Kowloon Tong, Hong Kong.*

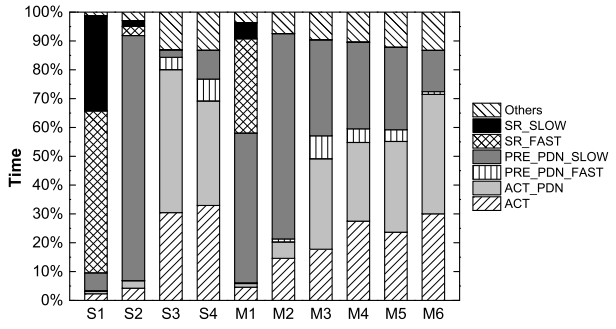

Fig. 2. The breakdown of time for RZ–SP.

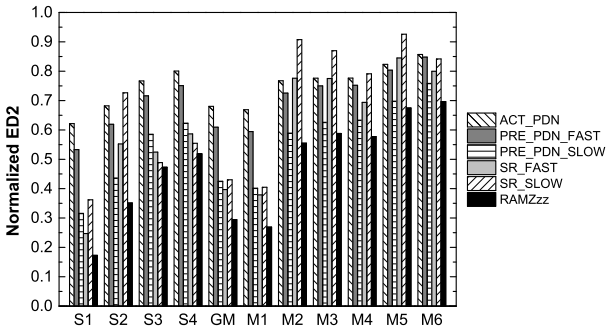

Fig. 3. Comparing ED^2 of RAMZzz and RZ–SD.

TABLE 1
Comparing ED^2 of RAMZzz with different number of low-power states on three DRAM architectures on M1.

Number of low -power states	1	2	3	4	5
DDR3	0.67	0.59	0.40	0.31	0.27
DDR2	0.68	0.43	0.35	0.30	N/A
LPDDR2	0.59	0.41	0.33	N/A	N/A

A.2 Studies on Adaptive Demotions

In this section, we study the impact of adaptive demotions, that is to compare the performance of RAMZzz and RZ–SD (RZ–SD uses the dynamic page migration without the adaptive demotion).

Figure 3 presents the comparison of ED^2 for RAMZzz and RZ–SD on DDR3. We compare the performance of RAMZzz with every possible RZ–SD approach on all workloads. That is, we use every available low-power state as the pre-selected low-power state in the RZ–SD approach. Since DDR3 has five low-power states, we have five RZ–SD approaches where each approach is denoted as the name of pre-selected low-power state Figure 3 (such as *SR_FAST* represents the RZ–SD approach which uses SR_FAST as the pre-selected low-power state).

We observe that RAMZzz outperforms all RZ–SD approaches on all workloads, with the reduction from 26.4% to 51.1% (36.4% on average). Moreover, different workloads have different choices on the most energy-efficient RZ–SD approach, indicating that the static demotion scheme can not adapt to different workloads. The efficiency of the static demotion scheme is closely related to the decision on the pre-selected low-power state, justifying the necessity of adaptive demotions. The total delay of RZ–SD is close to that of RAMZzz, less than 3% for all workloads.

Finally, we study the impact of the number of available low-power states. In Table 1, we change the number of available low-power states used on DDR3, DDR2 and LPDDR2 on M1 from 1 to 5, 1 to 4 and 1 to 3, respectively. We add a low-power state with smaller power consumption when increasing the number of available low-power states. As the number of available low-power states increasing, the normalized ED^2 becomes smaller. The improvement in normalized ED^2 by increasing the number of available low-power states from 1 to the maximum is 59.8%, 54.1% and 45.2% on DDR3, DDR2 and LPDDR2, respectively. This further proves the self-adapting feature brought by our proposed adaptive demotion scheme.

To summarize the impact of adaptive demotions, we observe RAMZzz has much lower ED^2 than RZ–SD on three DRAM architectures, and with the reduction of 26.4–51.1% (36.4% on average), 12.0–48.7% (25.0% on average) and 5.0–41.9% (22.4% on average) on DDR3, DDR2 and LPDDR2, respectively.