WPM P-3.12

A Dynamic-Voltage-Adjustment Mechanism in Reducing the Power Consumption of Flash Memory for Portable Devices

Li-Pin Chang, Tei-Wei Kuo, and Shi-Wu Lo

{d6526009,ktw,d89015}@csie.ntu.edu.tw

Department of Computer Science and Information Engineering
National Taiwan University, Taipei, Taiwan 106, ROC

ABSTRACT

A dynamic-voltage-adjustment method is proposed to reduce the power consumption of a flash memory storage system, depending on the system workload. The usefulness of the proposed method is demonstrated by a series of experiments with very encouraging results.

INTRODUCTION

With the increasing popularity of portable devices, there is a strong demand of power management mechanisms for microprocessors and other subsystems. Advanced microprocessor and flash memory chips design now supports several voltage levels and allows more intelligence in managing power consumption to lengthen the operating time of portable devices. The goal is to let a portable device operate at a low-voltage level, whenever possible, to save battery power. Recently researchers have started investigating voltage-clock-scaling scheduling algorithms for microprocessors, e.g., [1,2]. In particular, Douglass, et al. [4] provided a series of power consumption measurement for flash memory under different percentages of capacity utilization.

In this paper, we propose a dynamic-voltage adjustment method to manage the power consumption of flash memory. When the system is heavily loaded, a reasonable portion of operations will be performed at a high voltage level to deliver a reasonable performance. However, when the system workload is low, the proposed mechanism will try to service requests at a low voltage level without sacrificing the system performance. The strength of the proposed mechanism is evaluated by a series of experiments over a realistic trace and Intel NOR flash (see Table 1), for which we have very encouraging results.

SYSTEM MODEL

Because of the out-place update nature for flash memory, Flash memory Translation Layer (FTL) is introduced to emulate a block device for flash memory [3] so that users / applications have transparent access of data at dynamically

$V_{cc} = 5 \text{ v} \pm 10\%$, Block Size = 64KB

V_{DD}	Block write		Block erase	
	Power Consumption	Performance	Power Consumption	Performance
5v	375mW	0.5s	250mW	0.4s
12v	540mW	0.4s	480mW	0.3s

Table 1. Performance and power consumption of a typical NOR flash memory with different V_{pp} voltage supply [6]

allocated space. Note that available space is created by erasing and recycling blocks that are no longer used. Advanced flash memory design now supports multiple voltage levels. For example, the program / erase voltage supply V_{pp} can be either 5v or 12v, while read requests always use a fixed voltage supply V_{cc} . The characteristics of the typical flash we used is summarized in Table 1. For the purpose of our method, we assign a deadline to each request as a bound for its response time and to reflect the performance requirements of the system. Note that request deadlines are soft deadlines. Each request R_i is associated with an arrival time T_{a_i} , a completion time T_{f_i} and a deadline T_{d_i} , respectively. The time bet. T_{a_i} and T_{f_i} depends on the system workload and the supplied voltage V_{pp} .

ALGORITHM

Let $S = \{R_1, R_2, \dots, R_n\}$ be a collection of pending requests for flash memory. We propose to first schedule requests in S in an earliest deadline first (EDF) order [5]. Requests are

Figure 1. Concept of voltage-adjustment in scheduling

initially serviced at the low voltage supply. If there is any request that can not meet its deadline, we will raise the supply voltage of the first k requests until all requests meet their deadlines, or all requests are already serviced at the high voltage level. The number k must be minimized. The process in the voltage raising has a time complexity $O(n)$, where n is the number of pending requests. We shall illustrate the idea by an example, shown in Figure 1. Figure 1.(a) shows that R_2 will miss its deadline if all requests are serviced at the low voltage level. As shown in Figure 1.(b), when the first pending request, i.e., R_1 , is serviced at the high voltage level, all requests will meet their deadlines.

As astute readers may notice, the scheduling problem with minimal power consumption is inherently NP-hard because of the non-preemptivity of requests. In order to schedule requests efficiently, we propose an on-line algorithm with an $O(n)$ time complexity. The on-line algorithm is a greedy approach, which tries to satisfy a missed deadline by raising the voltage of the first k requests that are scheduled before the missed deadline, until the request deadline is satisfied or no further voltage raising is possible.

EXPERIMENT RESULTS

The capability of the proposed mechanism is evaluated over a typical NOR flash equipped with switchable Vpp voltage supply [6]. The workload traces were gathered by emulating web-surfing applications over a portable device. The primary performance metrics are the average response time of requests and the power consumption of writes/erases. The power consumption of reads was not considered because reads adopted a fixed voltage supply. Additional metrics such as the number of deadline violations and the average tardy time were also adopted.

The performance (in terms of average response time) and the power consumption of flash memory were evaluated under different system workloads, i.e., the mean arrival rate of requests. In this part, the write deadline was set as 200ms, and the read deadline was 10ms. It was shown that the average response time of requests under the proposed dynamic-voltage-adjustment algorithm was very close to that serviced at the high voltage level, as shown in Figure 2.(a), where "D", "H", and "L" denoted the results running at dynamic-adjusted, high, and low voltage levels, respectively. The power consumption of the system dropped rapidly as the system workload decreased (see Figure 2.(b)).

CONCLUSION

In this paper, we propose a dynamic-voltage-adjustment mechanism to fulfill the system performance and, at the same time, to minimize the power consumption. A scheduling framework is proposed and combined with a

voltage adjustment mechanism. The advantage of the proposed mechanism is demonstrated by a series of experiments over a typical NOR flash memory, for which we have very encouraging results.

REFERENCES

- [1] C.M. Krishna and Y.H. Lee, "Voltage-Clock-Scaling Adaptive Scheduling Techniques for Low Power in Hard Real-Time Systems," IEEE Real-Time Technology and Applications Symposium, 2000.
- [2] Y. Nakamoto, Y. Tsujino, and N. Tokura, "Real-Time Task Scheduling Algorithms for Maximum Utilization of Secondary Batteries in Portable Devices," International Conference on Real-Time Computing Systems and Applications, Dec. 2000.
- [3] A. Kawaguchi, S. Nishioka, and H. Motoda, "A Flash Memory based File System," Proceedings of the USENIX Technical Conference, 1995.
- [4] F. Douglass, R. Caceres, F. Kaashoek, K. Li, B. Marsh, and J.A. Tauber, "Storage Alternatives for Mobile Computers," Proceedings of the USENIX Operating System Design and Implementation, 1994.
- [5] C.L. Liu and J.W. Layland, "Scheduling Algorithms for Multiprogramming in a Hard Real Time Environment," Journal of the ACM, 1973.
- [6] "28F016S5 5-Volt FlashFile Flash Memory Datasheet", Intel Corporation, 1999

(b) Power Consumption of Write/Erase Requests

Figure 2. System performance and power consumption under different system workloads.