Current Trends in Flash Memory Technology

(Invited Paper)

Sang Lyul Min Eyee Hyun Nam

School of Computer Science and Engineering
Seoul National University
Seoul 151-742, KOREA
Tel : +82-2-880-7047
Fax : +82-2-886-7589
e-mail : symin@snu.ac.kr

Abstract - In this paper, we describe the basics of flash memory technology in general and flash memory drive in particular, and explain the current trends of major components of a flash memory drive including flash memory chips, host interface and flash memory controller.

I. Introduction

.

This paper explains the current trends of flash memory technology in general and flash memory drive in particular. The paper is organized as follows. Section II surveys the current trends of the two components that interface a flash memory drive: flash memory chips and host interface. Then, in Section III we explain various techniques to enhance the performance of a flash memory drive. Section IV briefly explains a recent storage system called a hybrid hard disk drive that combines the advantages of flash memory drive and hard disk drive. Finally, we offer conclusions in Section V.

II. Flash memory chips and host interface

As we can see in Fig. 1 that shows the basic architecture of a flash memory drive, there are two components that interface to the flash memory controller in a flash memory drive: flash memory chips and host interface.

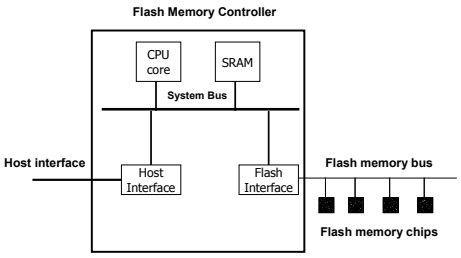

Fig. 1. Basic architecture of flash memory drive.

Fig. 2 gives a high-level interface of a NAND flash memory chip, the type of flash memory used for bulk storage, from a software perspective. The NAND flash memory chip consists of a set of blocks that in turn consist of a set of pages where each page has the data part that stores the user data and the spare part that stores meta data associated with user data such as ECC. Although different sizes may be used, currently the most popular block size is 128 Kbytes consisting of 64 pages of 2 Kbytes. There are three possible operations to a flash memory chip: **read page**, **program page**, and **erase block**. The read page operation, given the chip number, the block number, and the page number returns the contents of the addressed page, which takes about 20 us. Likewise, the program page operation writes the supplied contents to the target page and takes about 200 us. Unlike a write operation to other types of storage medium, the program operation can change the stored bits from 1 to 0 only. Therefore, the write operation is implemented by selectively changing bits from 1 to 0 to match the supplied contents assuming all bits in the target page are 1's before the program operation. In flash memory, the only way to change a bit in a page from 0 to 1 is to erase the block that contains the page. The erase operation sets all bits in the block to 1. The erase block operation takes about 2 ms.

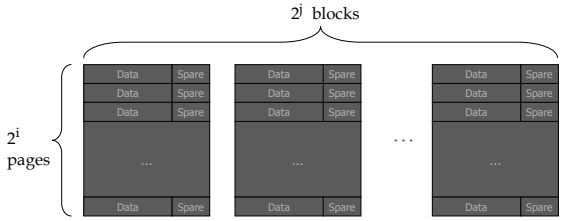

Fig. 2. NAND flash memory chip.

Fig. 3 shows the density and price trends of NAND flash memory chips. One notable point in the figure is the capacity of NAND flash memory has been doubled every year over the past five years. Another interesting point is that a flash drive whose capacity is comparable to that of hard disk drive used in a notebook these days (> 20 Gbytes) will become affordable (< $500) for high-end users within two or three years.

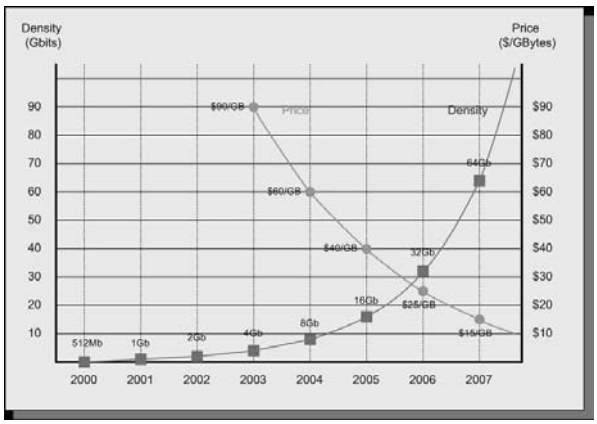

Fig. 3. NAND flash memory density and price trends.

0-7803-9451-8/06/$20.00

On the host interface side, the two most notable trends are (1) the maximum host interface bandwidth has increased exponentially and (2) most of recently proposed host interface standards are serial. We expect that these trends will continue for the foreseeable future and will assure a scalable performance for future flash memory drives.

III. Flash memory controller architecture

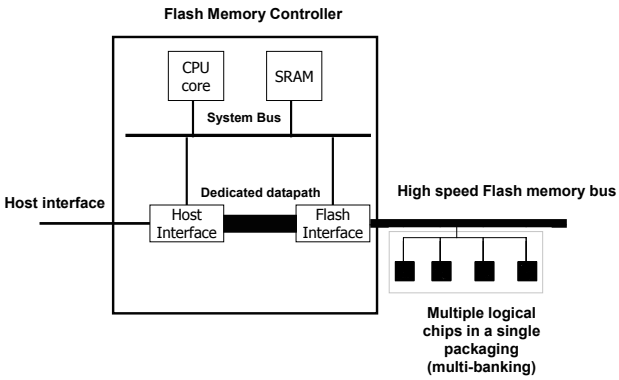

Fig. 4. Advanced flash memory controller architecture.

Fig. 4 shows various enhancements that can be made to the basic architecture of a flash memory drive in Fig.1. First of all, there is a lot of room for improvement in the architecture of flash memory chip itself. Currently, the maximum write bandwidth that can be provided by a single flash memory chip is 10 Mbytes/s (= 2 Kbytes/200us) that is certainly not sufficient for high-performance flash memory drives. Higher write bandwidth can be obtained by using interleaving within a chip or across multiple chips. This interleaving will give benefits to the read bandwidth and erase bandwidth as well.

Another place for improvement is flash memory bus. Currently the maximum flash memory bus bandwidth is 33 Mbytes/s with a 8-bit data bus. This means that the maximum performance of any flash memory drive that uses a single flash memory bus is limited to 33 Mbytes/s, which is also not sufficient for high-performance flash memory drives. Of course, multiple flash memory buses can be used for higher bus bandwidth but this approach will require a higher pin count and a bus interleaving logic within the flash memory controller. A preferred approach would be to use a single high-performance flash memory bus, an approach currently taken by the OneNAND bus architecture that provides up to 108 Mbytes/s with a 16-bit bus [1].

The third place for improvement is the datapath between host interface and flash memory. Traditionally, the datapath was implemented by a cascaded DMA transfers using SRAM within the controller but this approach not only increases the latency but also consumes a significant amount of system bus bandwidth making the system bus an overall system bottleneck that limits the maximum performance. An alternative approach is to use a dedicated datapath between host interface and flash memory as shown in Fig. 4 that not only reduces the latency but also has an advantage in terms of scalability.

IV. Hybrid hard disk drive

One of recent developments in storage systems is a so-called hybrid hard disk drive shown in Fig. 5. It adds a flash memory chip to a hard disk drive for reduced power consumption and faster start-up. In a hybrid hard disk drive, write requests from host are buffered in flash memory with the spindle motor spun down. When the flash write buffer is full, the spindle motor spins up and buffered requests are flushed to disk storage. This writing buffering using a 128 Mbytes of flash memory is shown to save up to 80% of power consumption in a 2.5in hard disk drive [2]. Moreover, by caching in flash memory those contents that will be used in the next boot/resume and using them while the hard disk drive is being spun up, the boot/resume time can be drastically reduced.

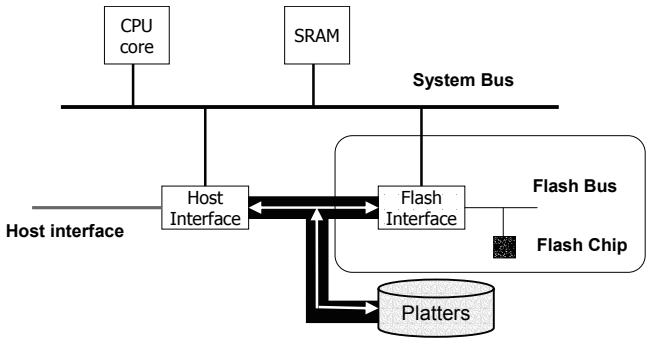

Fig. 5. Hybrid hard disk drive

V. Conclusions

This paper gave basics of flash memory drive and explained the current trends of its components including flash memory chips, host interface, and flash memory controller. The market of storage devices based on flash memory is expanding rapidly and this expansion will pose many challenges not only in flash memory technology but also in related areas as in the case of hybrid hard disk drives.

Acknowledgements

This work was supported in part by the IDEC (IC Design Education Center), by the IT-SoC project, and by Brain Korea 21 project. ICT at Seoul National University provided research facilities for this study.

References

[1] B. Kim, S. Cho, Y. Choi, and Y. Choi, "OneNAND(TM): A High Performance and Low Power Memory Solution for Code and Data Storage," Proc. 20th Non-Volatile Semiconductor Workshop, 2004.
[2] J. Creasey, “Hybrid Hard Drives with Non-Volatile Flash and Longhorn,” WinHEC 2005.