

Replacing HDDs with Solid-State Flash Disks in PXIbus-Based Systems

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

2006 J. Phys.: Conf. Ser. 48 1332

(<http://iopscience.iop.org/1742-6596/48/1/248>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 38.107.179.238

The article was downloaded on 26/05/2012 at 00:09

Please note that [terms and conditions apply](#).

Replacing HDDs with Solid-State Flash Disks in PXIbus-Based Systems

Z W Zhao and L Zeng

Academy of Equipment Command & Technology, Beijing 101416, China

Abstract. New security features, constantly decreasing prices, solid-state Flash disks are becoming a popular alternative for replacing failure-prone mechanical Hard Disk Drives (HDDs) in PXIbus-based military systems. The key component in high-capacity solid-state Flash disks is NAND Flash, but with a specification that shows only 100,000-300,000 write/erase cycles, engineers may be concerned that the lifetime of Flash disks cannot meet their application requirements. With the right Flash management, Flash disks are able to provide the reliability and endurance that military applications need.

1. In rugged environments the HDD can become the weakest link

The PCI eXtensions for Instrumentation (PXI) was created in response to the needs of a variety of instrumentation and automation users who require ever increasing performance, functionality, and reliability from compact rugged systems that are easy to integrate and use. Existing industry standards are leveraged by PXI to benefit from high component availability at lower costs. Most importantly, by maintaining software compatibility with industry-standard personal computers, PXI allows industrial customers to use the same software tools and environments with which they are familiar.

PXIbus systems were designed to effectively withstand shock, vibration, and extended temperature ranges, making them ideal for harsh environments and military applications. These same conditions, however, are causing HDDs to fail due to mechanical breakdowns such as read/write head failures and motor problems. With up to double-digit failure rates in rugged environments, the HDD can quickly become the weakest link in military systems. In addition, hard drives leave traces of confidential information after erasure.

2. Replacing HDDs with solid-state Flash disks

Flash disks are solid-state devices with no moving parts, able to operate in the harshest environmental conditions as defined in MIL-STD-810F: within -40 °C to +85 °C temperature range, absorbing shock conditions of an astounding 1500 G, and withstanding random vibrations of 16 G at 80,000 feet altitude. This makes Flash disks the ideal storage solution for rugged PXIbus-based military systems. In addition, the nature of NAND flash technology, whereby data is “burned” into the silicon to perform both read/write and erase operations, ensures that erased data cannot be retrieved. In fact, there is no indication of the number of erase cycles that have been performed from the time when a cell was programmed to zero.

In the past, cost was a real barrier to their deployment, but the flash industry has its own version of Moore’s law that overcomes this barrier: The density of the flash component doubles within the same silicon footprint size every 12 months. This enables double the capacity every year in the same casing size, while sharply reducing flash cost. Companies such as Vanguard Rugged Storage, Red Rock

Technologies, Asine, and Phoenix International provide reliable data storage by mounting Flash disks directly on 6U and 3U CompactPCI boards.

3. Flash technology limitations

All types of Flash memory have a specified life expectancy, measured as the number of erase cycles that each erase unit in the Flash device can undergo. One of the key ways to extend Flash life expectancy is by managing Flash. If Flash is not managed properly, its life expectancy can severely diminish the overall life of the Flash disk. On the other hand, using Flash management can substantially increase Flash life expectancy to the point where the limit is of no consequence in even the most write-intensive applications.

Solid-state Flash disks use NAND Flash technology as opposed to NOR Flash technology, due to its higher density (1 GB memory chip) and lower price per MB. But NAND Flash technology, in particular, is subject to spontaneous bit flips and bad blocks that can compromise data integrity. Here, too, Flash management can effectively overcome these inherent characteristics. Enhanced Error Detection and Correction Code (EDC/ECC) can provide a reliable NAND storage media that meets even the strictest requirements of military applications.

4. Flash management

There are several methods for managing the Flash media when using it to emulate a disk drive. Two of these methods are:

- The simple Flash algorithm method
- The wear-leveling method

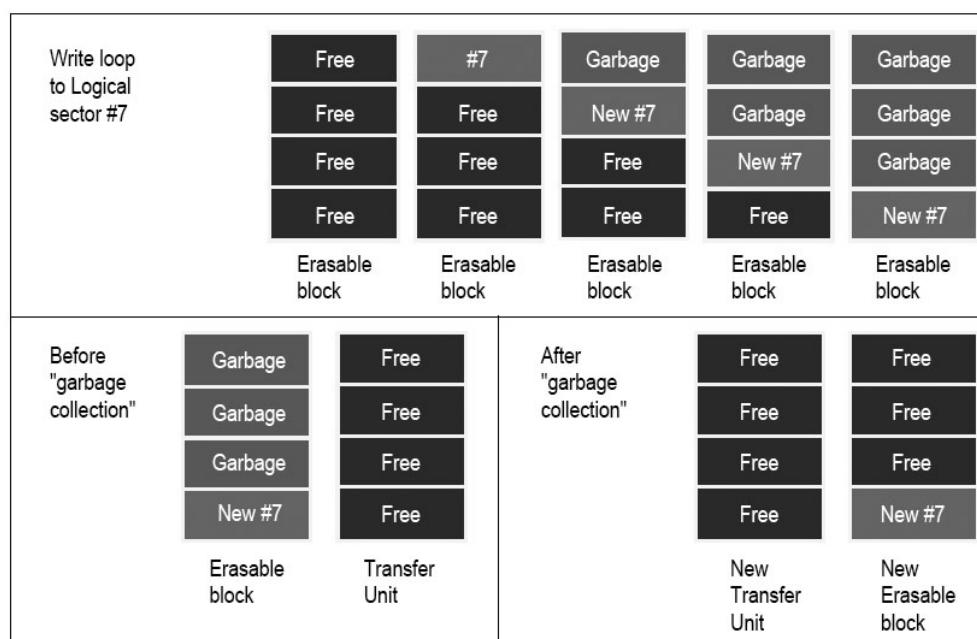


Figure 1. wear-leveling application in logical sector #7.

4.1. The simple Flash algorithm method.

Simple Flash algorithms map a logical sector to a fixed physical location on the Flash. This method quickly causes the Flash to wear out when an application updates the same sectors repeatedly. Updating the same group of sectors is a very common scenario, since all file systems need to maintain some data that describes the allocation of sectors to files, and this data is located in a specified area of a disk drive. For example, a File Allocation Table (FAT) file system, which updates the FAT every time a file is extended or concatenated, resides in sequential sectors located at the beginning of the

media. If these sectors are updated repeatedly, failures could result after only several thousands of file operations.

4.2. The wear-leveling method.

Advanced algorithms for flash management map a logical sector to a physical location that is not the same sector. This method, called wear-leveling, ensures that all write and erase cycles are evenly spread across the entire flash array to extend flash life expectancy. Static wear-leveling is applied on static files that are characterized by sectors of data that remain unchanged for very long periods of time. Dynamic wear-leveling is applied on newly written data based on a statistical allocation of units.

Figure 1 shows a wear-leveling application that is repeatedly updating data in logical sector #7. Flash management virtually maps logical sector #7 to a different physical location for every update. The new data is updated in different blocks, and the old data is marked as “garbage”, causing equal write/erase operations to all cells. Once an erasable block is full with garbage data, a garbage collection process is executed to refresh the block for a new write operation.

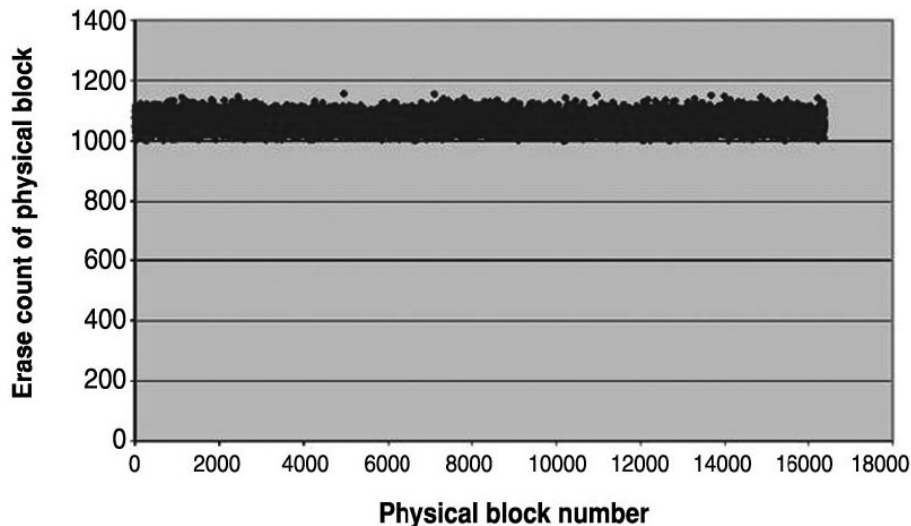


Figure 2. the media status after continuously updating the second logical half.

Figure 2 shows how wear-leveling extends Flash life expectancy by plotting the physical block erase count against the physical block number, effectively showing a “map” of the device’s blocks and how often they are erased. When implementing a wear-leveling scheme, all parts of the Flash are used in a statistically even manner, thereby extending the life expectancy.

The combination of both static and dynamic wear-leveling has the greatest effect on increasing flash life expectancy. More than extending flash life expectancy, wear-leveling also delays the onset of certain failure mechanisms in the flash. These failure mechanisms can cause entire erase units to become inoperable. When wearleveling is used, the erase cycle limit of the flash is increased beyond the minimum specified by flash vendors.

M-Systems’ True Flash File System (TrueFFS) Flash management tool, embedded inside its Fast Flash Disk (FFD) family of rugged Flash disks, incorporates one of the most effective wear-leveling mechanisms in the industry. Based on the performance of TrueFFS, Flash disks can guarantee more than 5 million write/erase cycles, compared with 100,000 to 300,000 write/erase cycles guaranteed by Flash chip vendors such as Toshiba and Samsung for raw NAND.

5. Life expectancy calculations

A number of methods can be used to calculate the life expectancy of the product in which the Flash is embedded, depending on the type of Flash management used. The most common Flash management algorithm used is called “Erase before Write.” Every time a sector (512 B) is written, a block of

sectors (16-128 KB) must be erased. The number of sectors in an erase block depends on the Flash chip used. When using this method, Flash disk endurance is calculated based on the number of sectors within an erase block, the media size (disk capacity), the write/erase cycles guaranteed by the Flash disk vendor (typically 100,000 to 300,000 cycles), and the rate at which data is written.

The calculations below compare the “Erase before Write” and TrueFFS technology on a 512 MB Flash disk with an application that fills the entire media 240 times a day (10 times an hour).

5.1. “Erase before Write” worst-case scenario calculation.

Each update to a sector (512 B) requires an erase of an entire erasable block (128 sectors, 64 KB). This scenario is typical for updating the File Allocation Table (FAT).

$$\text{Min Life Span} = \frac{100,000}{240 \times 128} \text{ days} = 3.25 \text{ days} = 0.0089 \text{ years}$$

5.2. “Erase before Write” best-case scenario calculation.

Disk updates are done once per full block (32 sectors). An entire erasable block needs to be erased for each update. This scenario is typical of data acquisition systems that update large files.

$$\text{Min Life Span} = \frac{100,000}{240} \text{ days} = 416.66 \text{ days} = 1.14 \text{ years}$$

5.3. TrueFFS technology best- and worst-case scenario calculation.

This scenario shows that wear-leveling provides a predictable life span, regardless of the usage scenario.

$$\text{Min Life Span} = \frac{5,000,000 \times 0.995}{240} \text{ days} = 20,729 \text{ days} = 56.79 \text{ years}$$

When using TrueFFS technology as the Flash management software, the worst-case Flash disk life expectancy is calculated based on the media size, the write/erase cycles guaranteed by M-Systems (committed to 5,000,000 cycles), the TrueFFS overhead factor (0.5 percent of the disk capacity used for TrueFFS internal overhead), and the rate at which data is written.

6. Conclusion

Reliable data storage is critical in PXIbus-based systems. The time, cost, and damage in reputation that may result from a faulty HDD can be devastating. Solid-state Flash disks can operate under the harshest environmental conditions to provide top data reliability. As the density of Flash continues to double every 12 months in the same silicon footprint, its cost will continue to decrease, making it more attractive as an HDD replacement.

The life expectancy and reliability of a Flash disk depends in large part on the Flash management used. Such as M-Systems’ TrueFFS Flash management, embedded inside its FFD family of rugged Flash disks, incorporates one of the most effective wear-leveling mechanisms to extend the life span of the Flash media. In addition, TrueFFS’ combined hardware and software EDC/ECC enables the company’s products to meet a reliability level based on the 5-nines (99.999%) concept.

References

- [1] PXI Systems Alliance 2004 *PXI Hardware Specification 2.2* 1
- [2] Ofer Tsur 2004 Flash Disks: Enabling Mission-Critical Systems *COTS Journal* 1 19
- [3] Samsung Electronics 2005 *1G x 8 Bit / 2G x 8 Bit NAND Flash Memory* 1.0 3
- [4] Esther Spanjer 2005 Flash management: How it extends the lifetime of solid-state Flash disks in VMEbus-based systems *VMEbus Systems* 8 26
- [5] Ofer Tsur 2003 Solid Security with Solid-State Flash Disks *Military Technology* 5 2