



WHITE PAPER

SANDISK FLASH MEMORY CARDS

WEAR LEVELING

October 2003

SanDisk® Corporation general policy does not recommend the use of its products in life support applications where in a failure or malfunction of the product may directly threaten life or injury. Per SanDisk Terms and Conditions of Sale, the user of SanDisk products in life support applications assumes all risk of such use and indemnifies SanDisk against all damages.

The information in this manual is subject to change without notice.

SanDisk Corporation shall not be liable for technical or editorial errors or omissions contained herein; nor for incidental or consequential damages resulting from the furnishing, performance, or use of this material.

All parts of the SanDisk documentation are protected by copyright law and all rights are reserved. This documentation may not, in whole or in part, be copied, photocopied, reproduced, translated, or reduced to any electronic medium or machine-readable form without prior consent, in writing, from SanDisk Corporation.

SanDisk and the SanDisk logo are registered trademarks of SanDisk Corporation.

Product names mentioned herein are for identification purposes only and may be trademarks and/or registered trademarks of their respective companies.

© 2003 SanDisk Corporation. All rights reserved.

SanDisk products are covered or licensed under one or more of the following U.S. Patent Nos. 5,070,032; 5,095,344; 5,168,465; 5,172,338; 5,198,380; 5,200,959; 5,268,318; 5,268,870; 5,272,669; 5,418,752; 5,602,987. Other U.S. and foreign patents awarded and pending.

Lit. No. 80-36-00278 10/03 Printed in U.S.A.

OVERVIEW

This purpose of this white paper is to help SanDisk customers understand the benefits of wear leveling and to assist customers in calculating life expectancy of SanDisk cards in specific applications.

Flash memory is susceptible to wear as a result of the repeated program and erase cycles that are inherent in typical data storage applications. Applications in which this is a major concern include hard disk replacement applications where write operations occur frequently. How a storage system manages the wear of the memory is key to understanding the extended reliability of the host that relies on these storage systems.

WEAR LEVELING METHODOLOGY

Current products available in the industrial channel use NAND flash memory. It is important to understand the NAND memory architecture to gain insight into the wear leveling mechanism.

Each memory chip is divided into blocks. A block is an array of memory cells organized as sectors. The number of blocks and sectors vary from product to product. The minimum unit for a write or read operation is a page (or sector). The minimum unit for an erase operation is a block. Physical blocks are logically grouped into zones. For the current technology, a typical zone size is 4 MB. However, this may change from product to product. Wear leveling is done within a zone. The current firmware does not spread the wear across the capacity of the card. Each zone has about 3% additional "spare blocks" beyond what is assigned to meet the logical capacity of the flash card. This group of blocks is commonly referred to as the "Erase Pool".

With the introduction of SanDisk's Write-before-Erase architecture, each time a host writes data to the same logical address (CHS or LBA), data is written into a newly assigned, empty physical block from the "Erase Pool". The intrinsic nature of writing to a new physical location each time a logical address is written to is the basis for wear leveling found in SanDisk cards. This action spreads the writes over the zone, thus greatly extending the overall life of the card. The methodology of using a large number of physical addresses to manage a smaller logical address table allows for rotation of the physical addresses among the entire group of physical blocks within a zone. The resulting wear leveling optimizes the effective life of the media and avoids prematurely reaching the end of life on frequently written to host addresses.

When a card detects that a block has reached the end of its useful life, it removes that block from the blocks that are available for write operations. The result is a reduction of the size of the erase pool. This does not affect the capacity of the card as seen by the host. When the pool of blocks available for write operations has been exhausted due to wear, the card will reach the end of its useful life for write operations.

Current SanDisk products do not preempt wear leveling events during normal operation of the card. Applications typically don't require such management beyond the natural wear leveling that occurs during normal host operations. As a result, the effectiveness of wear leveling in current SanDisk products is dependent upon host usage. It is important for customers whose applications do not fall into this typical usage pattern to understand how their applications will affect the lifetime of the card.

LIFE EXPECTANCY SCENARIOS

► best case analysis

In a typical application, large data files are written to the card occupying contiguous sequential logical address space. This results in optimal wear leveling and provides card life exceeding the specification for card endurance. This increased endurance is achieved as follows: The 2,000,000 endurance cycles specification (I-Grade only) is a result of large amounts of test data collected from a very large sample set that accounts for the extreme limits of the test population. With the 3% additional erase pool being used in an ideal fashion, the distribution is narrowed and the card will survive beyond its specified lifetime.

► worst case analysis

In the worst-case application, data will be written as single sectors to random addresses across the card. These single sector writes will exercise the erase pool more rapidly, requiring the system to perform a "garbage collection" operation to free up new blocks for subsequent write operations. At the extreme, each single sector write would cause one block to be programmed and erased. As a typical block size is 16kB or 32 sectors, the amount of wear is increased by a factor of 31 since 32 physical sectors are written and erased for each sector the host writes. Spreading this wear across the erase pool results in an effective 1/30 usable lifetime. This case is an extreme example and is only included to show the range of application dependence. This result is comparable to other vendor's cards based on memory with a 16kB erase block.

► analysis of host dependence

In assessing the life expectancy of a card in a given system several factors need to be understood. These factors include the types of files and their corresponding sizes, frequency of card write operations and file system behavior (including data structures). The types of files must be considered since some files, such as operating systems or executable files, typically remain in fixed locations once they are stored in the card. This limits the number of physical blocks available for circulation into the erase pool. The remaining capacity after these files have been accounted for can then be divided by the typical size of files that will be updated over the lifetime of the card. Related to this calculation is how the file system overwrites existing files. Typical operating system behavior, such as DOS, will allocate new blocks from the file allocation table, or FAT, and so repeated file writes will occupy a new set of addresses on the card. This is very beneficial in spreading wear across the card since it forces the card to cycle the entire physical

area being used for such files. Special cases to consider include those where the files being updated are very small. Typically an operating system uses a minimum number of sectors to store a file, referred to as a cluster. Typical cluster sizes range from 8 to 64 sectors in size. The cluster size is important for files that are the same or smaller than the 32-sector block since these may trigger garbage collection operations. If these updates happen in a random fashion (sequential updates would not be affected by cluster size) lifetime may be reduced as a result. Finally, the frequency of such updates is then used to determine how long it will take before the card reaches its statistical limit for endurance. These factors can be combined in an equation that can be used to calculate the minimum time a card will function in that application:

$$lifetime = 2,000,000 \times \frac{(C_{zone} - C_{fixed}) \times \left(1 - k_r \times \frac{32 - N_{cluster}}{32}\right)}{FS_{typ}} \times \frac{1}{f_w}$$

where C_{zone} is the total capacity of the zone, C_{fixed} is the capacity used by fixed files, $N_{cluster}$ is the cluster size, FS_{typ} is the average file size and f_w is the average frequency at which files are updated. k_r is a factor that is 0 for file sizes that are typically over 16kB or for applications that are not random in the order in which such files are updated.

Example 1

In this example 128 KB of data is updated once a day. The zone has 500 KB worth of fixed files. A 4 MB zone size is assumed.

$$lifetime = 2,000,000 \times \frac{(4000 - 500) \times (1 - 0)}{128} \times \frac{1}{1/day}$$

$$lifetime = 149828 \text{ years}$$

Example 2

This example is a data logging operation using a 1GB card where a 4kB file is updated every five seconds. This would result in sequential address being written.

$$lifetime = 2,000,000 \times \frac{4000}{4} \times \frac{1}{1/5 \text{ sec}}$$

$$lifetime = 317 \text{ years}$$

Example 3

This example is a data logging operation using the same 1GB card where a new 4kB file is written every five seconds. But in this case the cluster size is 4kB and it is expected that, due to file system fragmentation, the logical addresses will be written randomly.

$$lifetime = 2,000,000 \times \frac{4 \times \left(1 - 1 \times \frac{32-8}{32}\right)}{.004} \times \frac{1}{1/5 \text{ sec}}$$
$$lifetime = 79.3 \text{ years}$$

CONCLUSION

These examples are general in nature but show how the equation can be used as a guideline for calculating card lifetime in different applications. They also demonstrate that SanDisk card architecture exceeds reasonable life expectancy in typical applications. If a particular applications behaves in such a way that this equation cannot be applied, the SanDisk Applications Engineering group can assist in performing card lifetime analysis.

For more information, please visit the SanDisk Web site at: www.sandisk.com

SanDisk Corporation

Corporate Headquarters
140 Caspian Court
Sunnyvale, CA 94089
408-542-0500
FAX: 408-542-0503
URL: <http://www.sandisk.com>

SanDisk Corporation