An InnoDisk White Paper
August 2012

Wear Leveling Architecture And Read Disturbance Technology For Solid-State Drives

Prolonging SSD Lifespan
And preventing Read Disturbance

Revision History

Date	Version	Information
2012.07.09	1.0	First Release

Introduction

This white paper presents InnoDisk's wear leveling process for prolonging the lifespan of a solid-state drive (SSD) and preventing read disturbance.

Conventional file systems like FAT, UFS, HFS, ext 2 and NTFS were originally designed for magnetic strips and as such rewrite data repeatedly to the same area. However, SSDs can be put through a limited number of erase cycles before becoming unreliable. This is usually around 3,000/100,000(MLC/SLC) cycles. Wear leveling works around these limitations by spreading erasures and re-writes across multiple blocks in an SSD. In this way, no single block is subjected to constant read/writes.
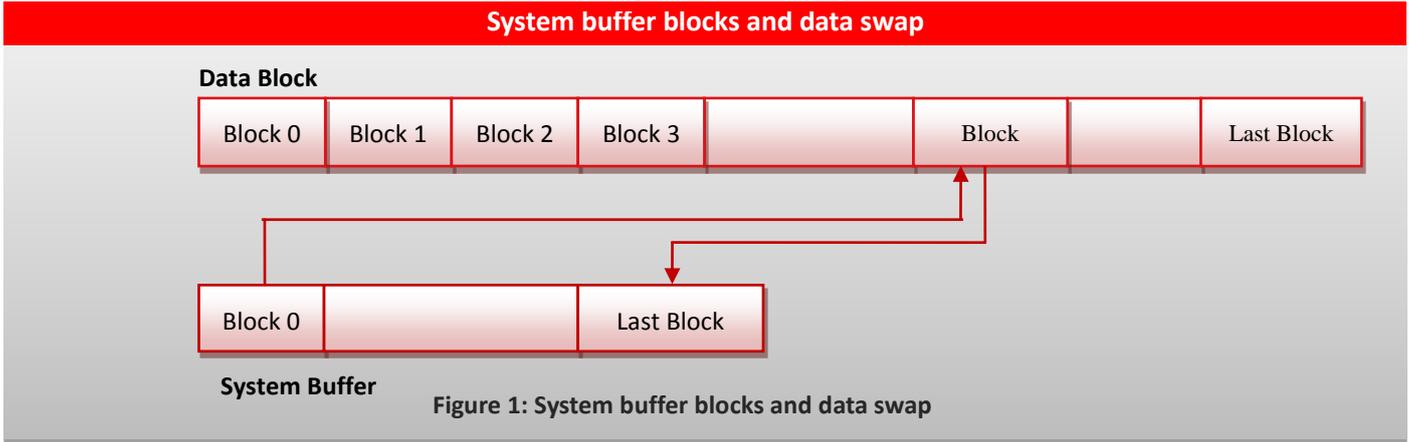
How InnoDisk's Wear Leveling architecture and technology prolongs SSD lifespan and prevents read disturbance

InnoDisk uses a combination of two types of wear leveling – dynamic and static wear leveling – to distribute write cycling across an SSD. Both types of wear leveling use a map to link Logical Block Address (LBAs) from the OS to the physical Flash memory. Static wear leveling, however, differs in that static blocks that do not change are periodically moved so that low usage cells are able to be used by other data.

System Buffer Blocks and Dynamic Wear Leveling

Because Flash blocks can't be overwritten, some spare blocks need to be reserved for temporary data buffering. These spare blocks can be used for other firmware purposes. In this white paper, these spare blocks are refered to as "system buffer blocks".

The simplest way to perform wear leveling is to recycle the programmed block (Block N in Figure 1) with the blocks in the system buffer block list. See Figure 1. If the host wants to overwrite Block N, the 1st block in the system buffer block would be used to replace Block N. Block N would be erased and put into the system buffer block list.

System buffer blocks and data swap

Data Block

Block 0	Block 1	Block 2	Block 3		Block		Last Block

Block 0		Last Block

System Buffer

Figure 1: System buffer blocks and data swap

In the worst case scenario the host keeps writing data to Block N while the system buffer blocks are used frequently and the other data blocks are not used at all. All system buffer blocks would be worn out after around
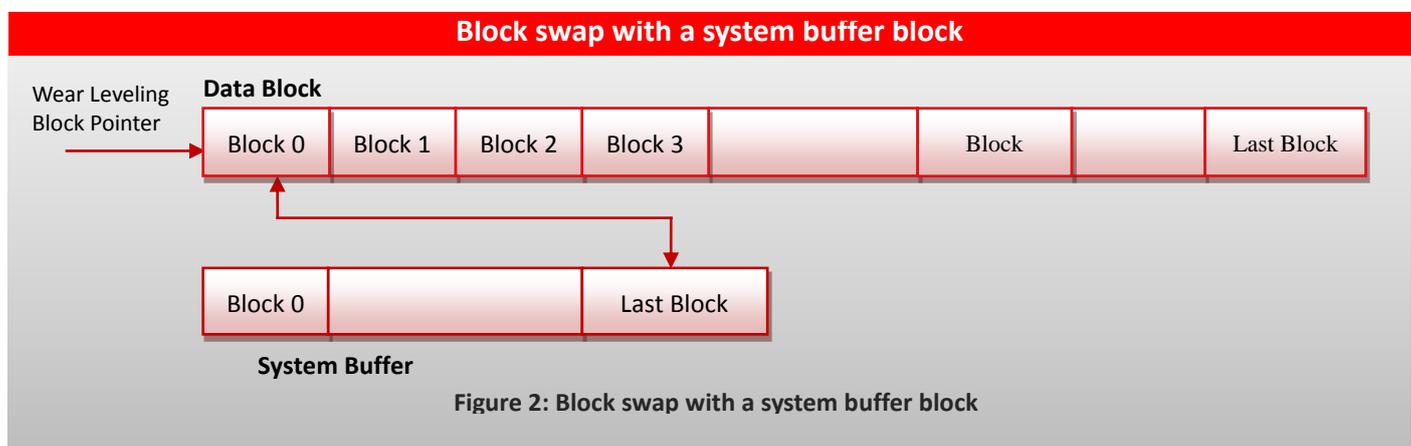
$$((\text{\# of system buffer blocks} + 1) * \text{Endurance})$$

This type of process is referred to as Dynamic Wear Leveling. It is not sufficient to only use dynamic wear leveling for solid-state drives as some operating systems run continuous background data swaps. The erase count would increase significantly in a short period of time. Static wear leveling must be added to the architecture.
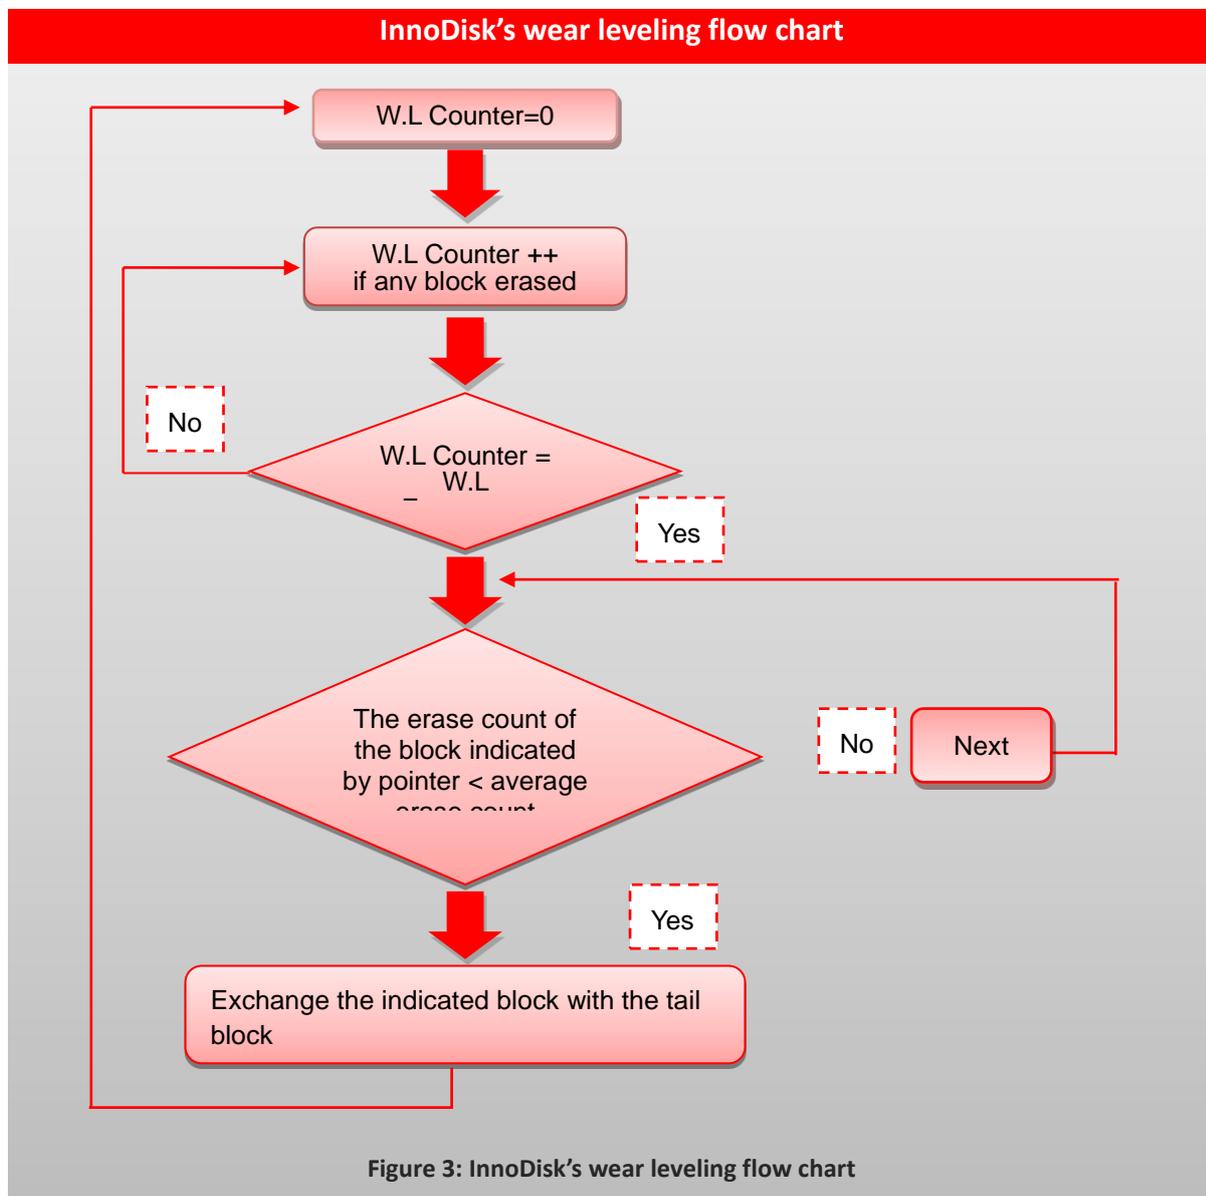
Static Wear Leveling

An innovative technology is adopted in InnoDisk's static wear leveling process to spread the program/erase count evenly on all blocks. Wear leveling performance frequency can be configured by customers in the .INI file. The variable name is *WearLevelFrequency*.

InnoDisk's wear leveling is run in the background. The internal micro controller increases a wear leveling counter whenever a block-erase is done. If the counter hits the defined wear leveling frequency, the controller initiates block swapping. See Figure 2.

Block swap with a system buffer block

Wear Leveling
Block Pointer

Data Block

Block 0	Block 1	Block 2	Block 3		Block		Last Block

Block 0		Last Block

System Buffer

Figure 2: Block swap with a system buffer block

The controller checks if the erase count on the block indicated by the pointer is less than the average erase count. If it is less, it means that the block is a less frequently used block. It is therefore suitable to swap the block with a system buffer block as system buffer blocks are generally more frequently used. If the erase count of the checked block is higher than average then the controller will move the pointer and check the next block until a block is found. See Figure 3 for the InnoDisk's wear leveling flow chart.

InnoDisk's wear leveling flow chart

W.L Counter=0

W.L Counter ++
if anv block erased

No

W.L Counter =
W.L

Yes

The erase count of
the block indicated
by pointer < average
erase count

No   Next

Yes

Exchange the indicated block with the tail block

Figure 3: InnoDisk's wear leveling flow chart

Read Disturbance

"Read Disturbance" is an anomaly of Flash IC, which occurs when one or more bits are changed during a READ operation. Read disturb errors occur within the block being read, but on a page or pages other than the page being read. Performing very high number (thousands or millions) of READ access on individual pages before an ERASE command for the block containing those pages can exacerbate this error.

Solution

InnoDisk Firmware is designed to resolve this issue by wear leveling and refresh ("re-charges").

With wear leveling feature, not only spread the program/erase count evenly on all blocks, but also can reduce the read access frequency to prevent Read Disturbance by reprogram the data to different block. Alternatively, ECC (Error Correcting Code) can detect and fix the data where the electrical properties may have been altered by refresh. When error bits in a block reach a threshold of say 17 error bits out of 24 bits, the block is automatically refreshed. i.e. the data is deleted and re-written. This stops the controller from constantly reading blocks with too many error bits and prevents read disturbance.

Conclusion

InnoDisk's wear leveling architecture and technology delivers an innovative method of spreading write/erase operations on an SSD evenly across blocks. By using system buffer blocks along with sophisticated block swapping and a configurable wear leveling performance frequency setting, SSDs using this technology can prolong their lifespan and prevent read disturbance.

InnoDisk SSD series at InnoDisk support wear leveling architecture and solution for read disturbance

2.5''&1.8''Industrial SSD Series, SATADOM Series, mSATA &mDOM-U Series, SATA Slim Series, CF-SATA & CFast Series, CF Card, SD, mSD & USB Series

About us

InnoDisk is a worldwide leading provider of data storage and memory module solutions for industrial and mission-critical applications. Leveraging in-house engineering and R&D expertise with a keen insight on industry trends, InnoDisk's solid-state drive (SSD) technologies provide enhanced, vertically-integrated data storage solutions. Our advanced Flash-based data storage and DRAM memory solutions meet stringent aerospace and defense application requirements, and are also widely used in industrial applications and embedded systems. InnoDisk offers customized solutions, from unique form factors to special firmware designs, and our support team of hardware, software and firmware engineers is always ready to tailor the right solution to each customer's needs. InnoDisk continually strives to innovate and provide system integrators and end customers with the best service in the industry.

For more information on InnoDisk's product line, technologies and applications, please visit

www.innodisk.com