

White Paper

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



Introduction to Wear-Leveling: Ensuring the Longevity of SSD

This white paper presents Innodisk's Wear-Leveling process to ensure the longevity of solid-state drives (SSDs) and prevent read disturb and data retention issues.

Traditional file systems such as File Allocation Table (FAT), Unix File System (UFS), Hierarchical File System (HFS), extended file system2 (ext2), and New Technology File System (NTFS) were initially designed for tape, causing data to be repeatedly rewritten to the same area. However, SSD becomes unreliable after a limited number of program/erase cycles (P/E cycles) (around 3K/3K/60K for 3D TLC/MLC/SLC).

When too many P/E cycles cause the number of unreliable blocks to approach the size of OP (over-provisioning), the SSD switches to read-only mode to protect the stored data. This can mean the end of the SSD's life for users. However, at this point, there may still be blocks with fewer P/E cycles that are available.

Wear-Leveling balances the P/E cycles of all blocks in SSD to avoid premature termination of SSD life.

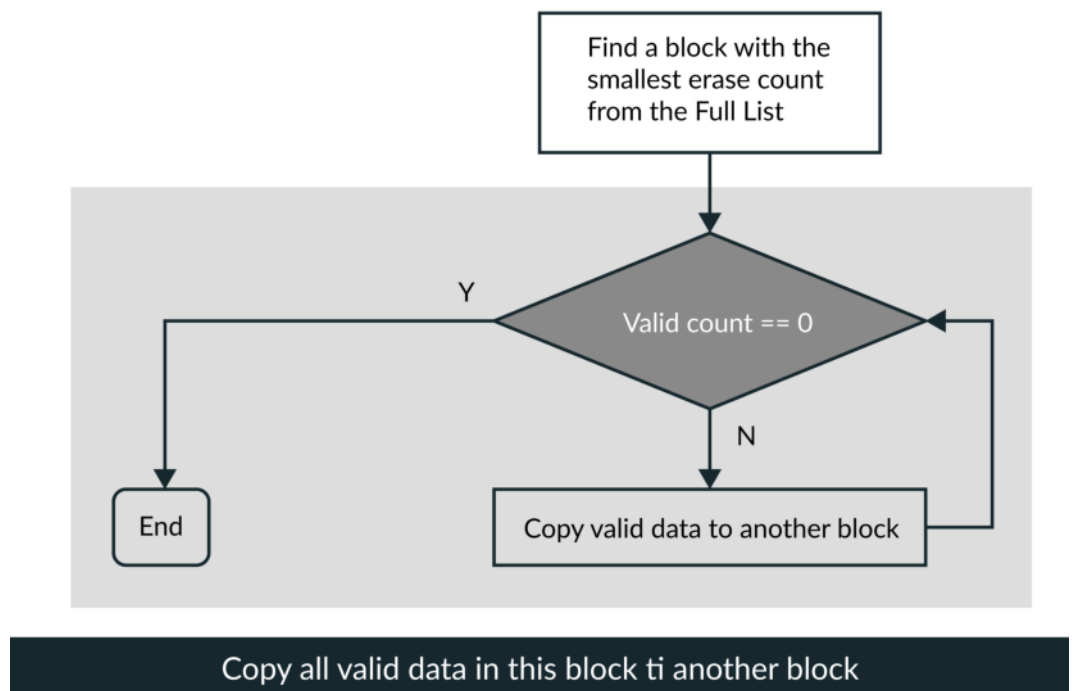
How Innodisk's Wear-Leveling Architecture and Technology Ensures the Longevity of SSDs and Prevents Read Disturb

Innodisk's Wear-Leveling mechanism combines dynamic and static mechanisms.

- **Dynamic Wear-Leveling:**
Dynamically monitor the balance status of all block P/E cycles. When the balance status exceeds the threshold, Dynamic Wear-Leveling is initiated to allow low-use blocks to write data.
- **Static Wear-Leveling:**
Regularly move unchanging static blocks in sequence, allowing low-use blocks to write data. This can also prevent read-disturb and reduce data retention issues.

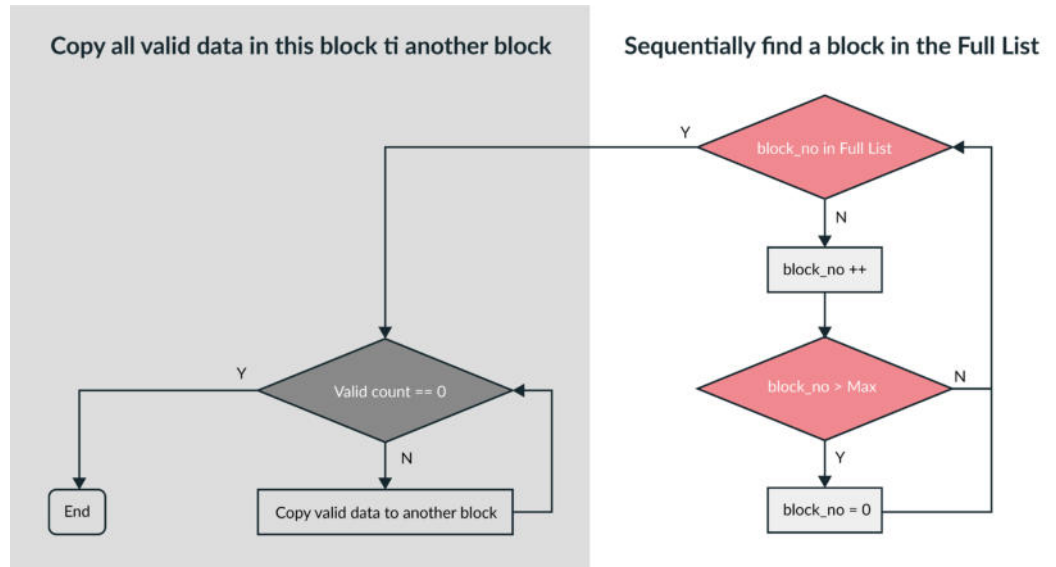
Dynamic Wear-Leveling

1. Find the block with the most minor P/E cycle counts in the SSD.
2. Copy the valid data in this block to another block with a larger P/E cycle count.
3. When no valid data is left in this block, it can be released for future usage.



Static Wear-Leveling

1. Periodically copy the valid data in static blocks to another block with larger P/E cycle counts.
2. When no valid data is left in this static block, it can be released for future usage.



Read Disturb

“Read disturb” occurs when data is read repeatedly from a particular memory cell or block. This can cause electrical charges to build up in adjacent cells, leading to errors in reading and writing data.

Over time, read disturb can lead to data corruption and reduce the overall lifespan and reliability of an SSD. To mitigate these effects, wear leveling plays an important role.

Data Retention

SSDs require power to maintain their data. If left unpowered for too long, the electrical charges in the NAND flash memory that stores the data may leak, causing data loss or corruption, generally referred to as “Data Retention.”

Innodisk Solution

Innodisk’s firmware solves this problem through Wear-Leveling. Wear-Leveling not only evenly distributes erase times to all blocks but also reduces the read frequency of each block to prevent read disturb and reduce data retention issues.

Conclusion

Innodisk's Wear-Leveling architecture and technology provide an innovative way to distribute write/erase operations on SSD evenly. By using system buffer blocks, complex block exchanges, and configurable settings for wear balance performance frequency, SSDs utilizing this technology can ensure their longevity and prevent read disturb and data retention.

Innodisk Corporation

5F., NO. 237, Sec. 1, Datong Rd., Xizhi Dist., New Taipei City, 221, Taiwan

Tel : +886-2-7703-3000

Fax : +886-2-7703-3555

E-Mail : sales@Innodisk.com



innodisk

Copyright © October 2023 Innodisk Corporation. All rights reserved. Innodisk is a trademark of Innodisk Corporation, registered in the United States and other countries. Other brand names mentioned herein are for identification purposes only and may be the trademarks of their respective owner(s).