NAND Flash Management Technology for Performance Enhancement

Author: Ares Cheng

Keyword: SSD, Over-Provisioning, NAND Boost

Performance, reliability, and endurance are key features of Solid State Drives (SSD) as they gradually replace HDDs as storage device of choice for commercial and industrial applications. Most SSD based storage uses NAND flash which is a non-volatile and random access type of memory.

Industrial-grade NAND flash storage devices add extra layers of reliability and performance required for today's mission critical applications. From SLC, to 3D NAND, BiCS 3, the unit storage density has grown rapidly but the performance, reliability, and endurance have decreased. So it's more important than ever to develop controller firmware to properly manage NAND flash to help improve its performance and reliability.

The purpose of this paper is to introduce how Advantech SQFlash develops NAND flash management technologies to enhance storage devices to meet the demands for high capacity, performance, and industrial grade reliability. This includes the introduction of Over Provisioning (OP) technology for improving performance and reliability, and NAND Boost technology for data transfer speed improvements.

Table of Content

Over-Provisioning

For endurance and performance enhancement, there are several widely used technologies for NAND flash management, such as garbage collection (GC), wear-leveling, and multi-write technology. Over-provisioning (OP) is an advanced technology that delivers better reliability, higher sustained performance, and extended endurance. Over-provisioning is the difference between the physical capacity of flash memory and the logical capacity presented by the operating system (OS) as available to the user. In other words, this OP capacity is non-user accessible and invisible to the host operating system. It is strictly reserved for the SSD controller's use.

Over-Provisioning Setting for Applications

The equation for calculating the percentage of over-provisioning is as below:

$$\text{Over Provisioning \%} = \frac{\text{Physical Capacity} - \text{User Capacity}}{\text{User Capacity}} \times 100\%$$

For instance, a total physical configuration consisting of 256 Gigabytes (GB) of flash memory has only 240 GB available to the user and the OP percentage is 6.67%, which is typically rounded up to 7 percent:

$$6.67\ \% = \frac{256\ GB - 240\ GB}{240\ GB} \times 100\%$$

There are many applications in the IT industry. For individual applications, OP percentage values are also individual suggested values. As shown in Table 1, take a 256GB SSD for example, the OP of consumer electronics applications (clients) is commonly set for 7%. For enterprise-oriented applications that focus on reliability and endurance, OP is usually set to around 28%, and in some high-end enterprise and industrial applications, OP is set at 50%.

Physical Capacity	User Capacity	Over-Provisioning %	Application Model
256 GB	128 GB	50%	Enterprise (Write Intensive)
256 GB	200 GB	28%	Enterprise
256 GB	240 GB	7%	Client

Table 1: Over-Provisioning Applications

In addition to enhancing the efficiency of garbage collection, wear-leveling, and bad block management, OP has three major benefits:

Better Sustained Performance

By utilizing the additional space from OP, SSD controllers can enhance the efficiency of garbage collection and their sustained performance improves. 4K random write tests with different OP percentages were tested and the results shown in Figure 1. The colored lines show how 4K random write performance (IOPS) will be affected by how much OP percentage is applied. Obviously, IOPS is significantly improved by increasing the OP percentage.

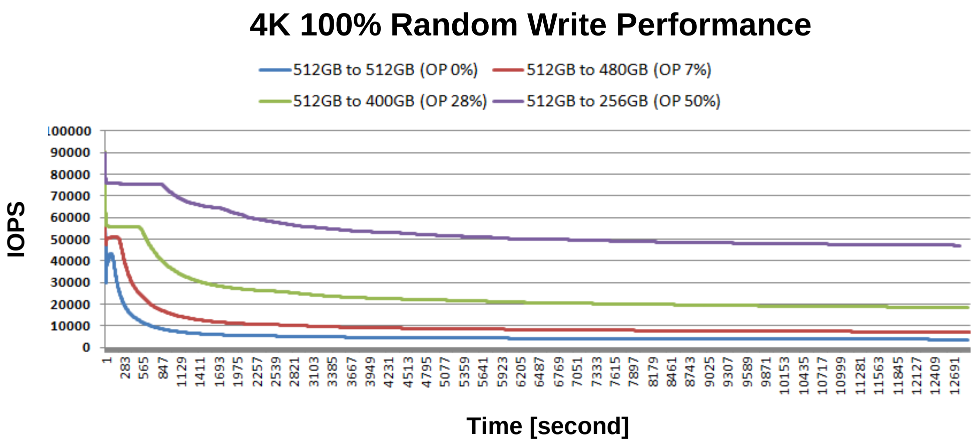

Figure 1: IOPS vs. Over-Provisioning

Smaller Write Amplification

A host system writing data into NAND flash memory generally includes moving valid data during multiple garbage collection cycles. Therefore, the total written data size to NAND flash is larger than the issued data size from the host system. For example, if a host system only issues 1GB data but 2GB data is written into NAND flash, this phenomenon indicates a write amplification (WA) of 2.0. And if the WA of the NAND flash memory is too big, it might seriously impact performance and endurance. Over-provisioning gives a positive influence on the write amplification. By increasing the OP percentage, the NAND flash controller, along with flash management algorithms can run more efficiently. The correlation between WA and OP is shown in Figure 2 below. This experimental result was tested by the JEDEC 219 client workload.

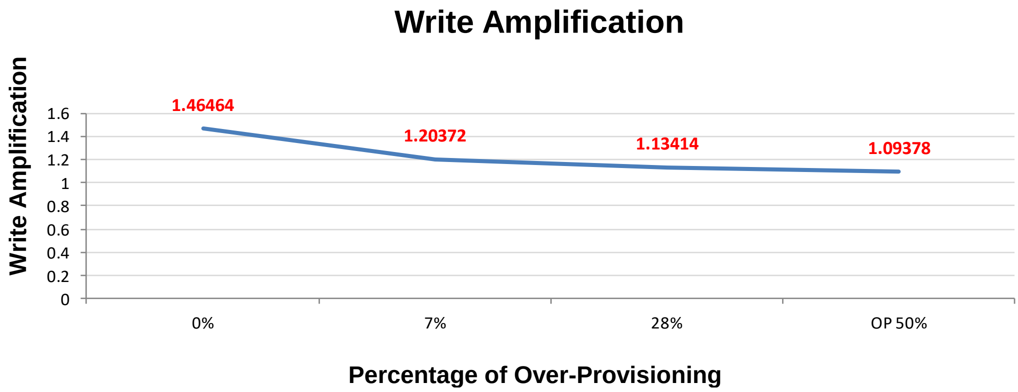

Figure 2: Write Amplification vs. Over-Provisioning

The test results indicate that the higher the OP percentage, the more WA can be decreased and the higher efficiency of garbage collection.

Extended Endurance

To estimate the endurance of NAND flash memory devices, terabytes written (TBW) is usually used as a reference to represent NAND flash P/E cycles. TBW is the estimated total amount of data that can be written into NAND flash memory before it reaches its lifespan. The TBW formula is defined by JEDEC, and write amplification is one of the key factors. With larger OP percentages, SSD write amplification is smaller as Figure 2 shows, and achieves larger TBW. Therefore, by increasing the OP percentage, the endurance of NAND flash memory is improved.

$$Terabytes\ Written\ (TBW) = \frac{PE\ Cycle \times Drive\ Capacity\ (GB)}{Write\ Amplification\ \ x\ 1000}$$

NAND Boost Technology

Introduction of SQFlash NAND Boost

When the storage bit of each cell is increased, the read and write performance of the NAND flash is reduced accordingly. This issue is often criticized by users. To solve the problem of low read and write performance of BiCS 3 or planar TLC SSDs, the SSD controller divides a certain proportion of space according to the capacity of the SSD to simulate SLC NAND for a performance boost. Since SLC stores only 1 bit of data per cell, the read and write performance is significantly ahead of planar TLC and BiCS 3. However, special attention should be paid to the fact that the NAND Boost area has to encounter very frequent data read and writes of internal data and table management. For some cases, the SSD controller divides an area of specific physical block addresses (PBA) in the NAND flash for simulating SLC. However, due to this kind of working model, the P/E cycle of the PBA will be exhausted more quickly and bad sectors will be more easily generated.

In order to avoid the issue of quickly consuming P/E cycles, Advantech SQFlash team developed a different working model for boosting NAND flash that uses the entire space of the SSD to dynamically divide the simulated SLC area for the boost function. That is to say, the NAND boost area is managed by the logical block address (LBA), and it is dynamically re-mapped with proper PBA according to the advanced wear-leveling algorithm that SQFlash adopted. Also, the size of the NAND boost is variable depending on the actual capacity of the SSD by a certain percentage. This design can not only improve performance, but also maintain the endurance of SSD.

Performance Enhancement Verification

In order to verify the relative performance enhancement of NAND Boost, benchmarks were set with NAND Boost function enabled and NAND Boost function disabled for performance tests on the same SSD. Test conditions were to continuously perform sequential write tests in one minute. The results are shown in Figure 3. Performance with NAND Boost enabled was around 500MB/s, which is about 2.5 times better than with NAND Boost disabled.

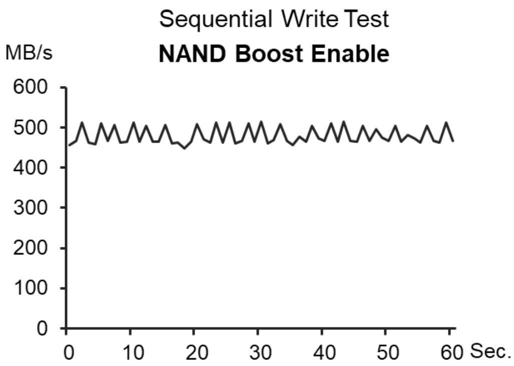

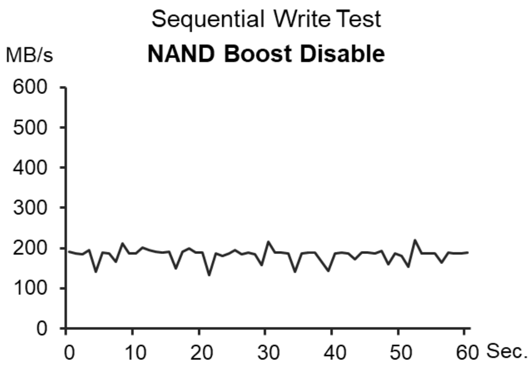

Figure 3: Sequential Write Tests of NAND Boost Enabled and Disabled

To better understand the sustained performance of NAND Boost, another benchmark experiment was designed to extend the time to three minutes of sequential write tests. Results from the experiments showed that the performance drastically decreased after 65 seconds from an average of 500 MB/s, to an average of 200 MB/s as shown in Figure 4. Therefore, after the SLC space gets used up, performance will quickly decrease to the TLC level.

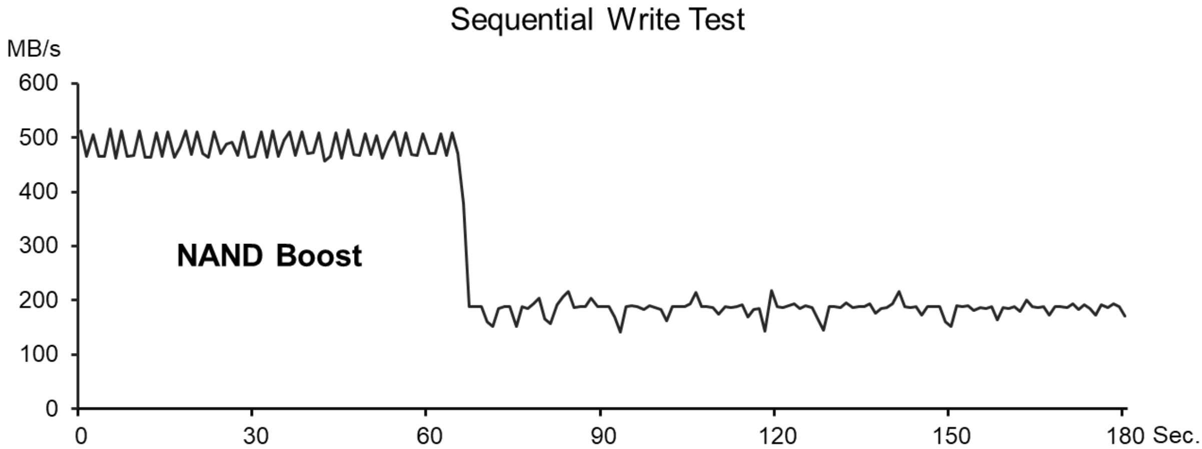

Figure 4: Long-term Sequential Write Test of NAND Boost

To verify the recoverability of NAND Boost, this experiment was designed as a long-term test. Test conditions were: After a long period of sequential writes, performance stabilized at the TLC level, the system then entered an idle state for a period of time, then the test started again and repeated. The experimental results are shown in Figure 5. After a period of idle time, the SLC space can be reconstructed again by the SSD firmware algorithm and the NAND Boost function performed again.

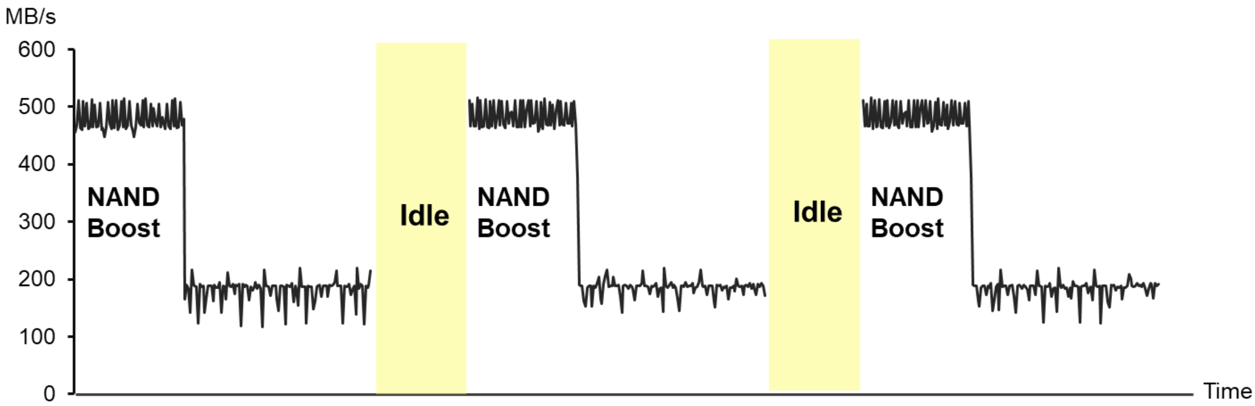

Figure 5: Long-term Sequential Write Test of NAND Boost

Benefits and Application of NAND Boost Technology

According to experimental results, NAND boost technology can greatly accelerate data transmission rates over a short time, but due to the limited size of NAND boost, it cannot provide sustained acceleration. This effect is very useful, especially for the SSD devices which use entry level SSD controllers of low-power consumption with TLC or BiCS 3, the speed of data transmission can be magnificently boosted. This kind of benefit is very suitable for systems which need fast boot up or fast shut down, and intermittent reading and writing of data. For example, POS machines, ATMs, and Kiosks, and industrial tablet devices are all suitable applications. Additionally, Advantech SQFlash team can customize the NAND boost space configuration based on the customer's requirements.

Conclusion

NAND flash technologies continue to evolve and bring higher storage densities, but they may also sacrifice the endurance, reliability, and data transfer speeds that are important for industrial storage. In order to consolidate the performance of NAND flash storage devices, the Advantech SQFlash team continues to research and develop NAND flash management technology. The Over-provisioning technology mentioned in this paper can be further combined with technologies such as Garbage Collection, Wear-Leveling, and Multi-Write functions to achieve the advantages of better sustained performance, smaller write amplification, and extended endurance. For data transfer speed improvements, NAND Boost technology can increase the write speeds of TLC level to SLC levels in a short period of time. SQFlash NAND Boost Technology uses dynamic methods so that the endurance can be maintained. In addition, Advantech SQFlash offers customization services according to the customer's application requirements.

For more information about industrial storage modules SQFlash, please visit sqflash.advantech.com