

How Micron SSDs Handle Unexpected Power Loss

The Problem

Multilevel cell (MLC) NAND technology has enabled high-performance data storage to be widely available at affordable price points—but as with all technologies, MLC NAND has its limitations. For example, all MLC devices are vulnerable to data loss in the event of an unexpected power loss. This paper discusses the data protection schemes employed by Micron’s client and enterprise solid state drives (SSDs) to eliminate unexpected power loss as a concern for customers and end users. Due to the differences in power-loss protection between client and enterprise drives, it is important for end users and system administrators to understand how a particular drive will respond in their application.

Handling an Expected Power Loss

The source power to SSDs is shut down on a frequent basis. Under normal circumstances, when a computer system is shut down (or when it goes into sleep or hibernate mode), the host computer sends a signal to the storage device to indicate that the power is going to turn off. In SATA systems, a STANDBY IMMEDIATE command is commonly used; in SAS or SCSI systems, the proper command is included in the start-stop protocol. The host system does not shut down until the command is sent and the drive acknowledges the command back to the host computer. The SSD completes any unfinished writes and saves any updated data addressing information before sending this acknowledge message.

What is an Unexpected Power Loss?

An unexpected power loss, on the other hand, is a power loss that is not preceded by a shutdown

notification; therefore, the drive cannot close out the final operations, which can lead to data and addressing information being lost. Other common terms to describe this situation are “surprise power loss,” “dirty power loss,” and “asynchronous power loss.”

How MLC NAND Works

Most MLC NAND available in the industry—including Micron’s—has two bits per cell. (NAND with three bits per cell is typically known as three-level cell [TLC] NAND, which has similar issues, but will not be discussed in this paper.) The NAND storage cells are arranged in pages that each typically contain 4KB to 16KB of capacity. One page is the smallest unit of memory that can be written in one operation. Typically, entire pages are programmed simultaneously and arranged in lower and upper pages. Because the lower pages are easier and faster to program, Micron’s Flash translation layer (FTL)—the translation table between the logical block address (LBA) and the physical location—favors lower pages when programming new data from the host computer. The lower-page program is relatively simple, as shown in Figure 1.

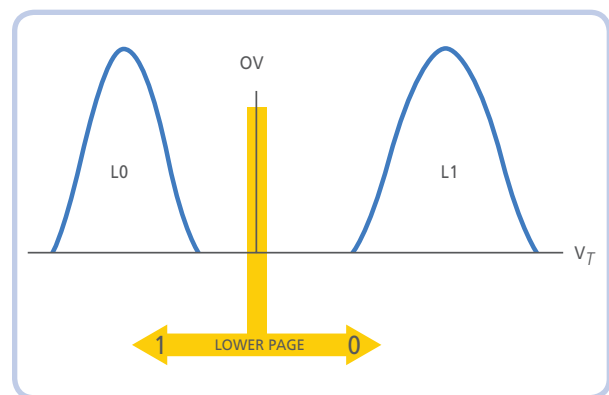


Figure 1: Possible MLC Signal Distributions – Only Lower Page Written

Programming NAND cells is an exercise in signal distributions. As shown in Figure 1, V_T is the threshold voltage that is used to detect data bits in the cell. In this case, if the stored signal rests in the L0 distribution (a negative V_T), the data bit is considered to be a value of 1 while a voltage detected in L1 (a positive V_T) is considered to be a value of 0. The scenario is complicated when the adjacent upper page is programmed, as shown in Figure 2.

If the lower page is programmed to the L0 like in Figure 1, the upper-page program moves the signal to L0 or L1, as shown in Figure 2. The new signal is

then decoded as 11 at L0 or 10 at L1. If the lower page is programmed to L1, as shown in Figure 1, the upper-page program moves the signal to 00 at L2 or 01 at L3, as shown in Figure 2.

The complication arises when the upper-page program is interrupted while this signal is being moved up in voltage (to the right in Figure 2)—like during an unexpected power loss. The incomplete program could end up incorrectly at any level or indeterminately in between the detection levels, which may result in an unfinished upper-page program, or worse, corrupted adjacent lower-page data that may have been programmed in the distant past.

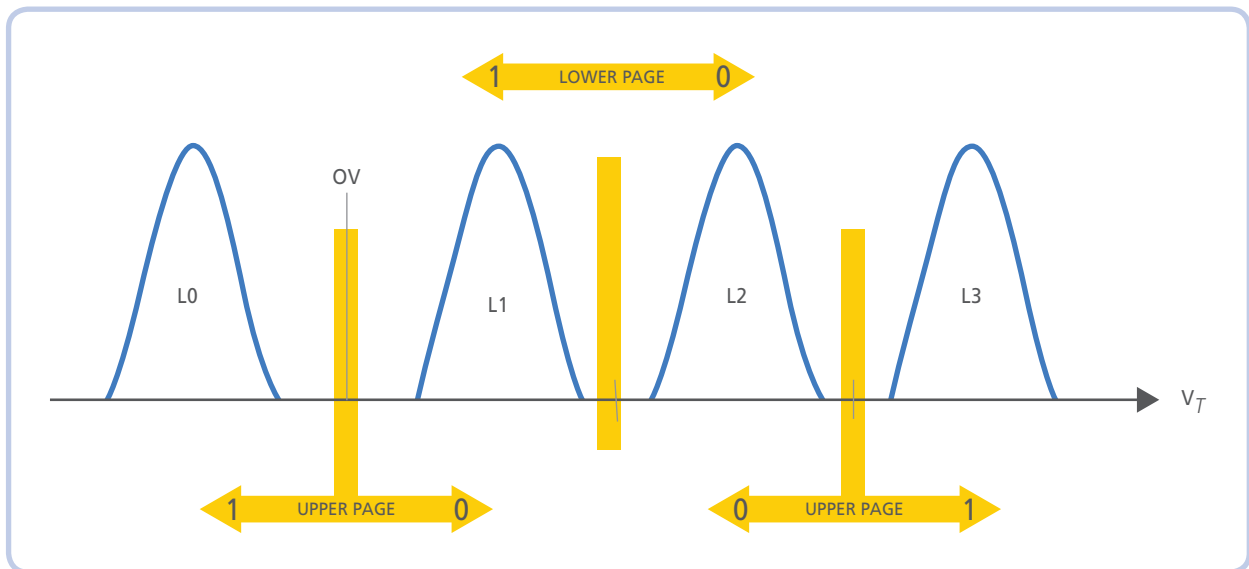


Figure 2: Possible MLC Signal Distributions – Upper and Lower Pages Written

Countermeasures in Client Computing

In client computing (desktop, notebooks, and professional tablets), “data at rest” refers to host or user data that has been physically saved to the storage media. “Data in flight” refers to writes sent to the drive from the host computer but not yet committed to the media—or any write that is in progress, but is not complete. Data in flight also refers to data that has been temporarily buffered or cached in volatile memory, acknowledged as written back to the host, but not yet committed to the NAND device.

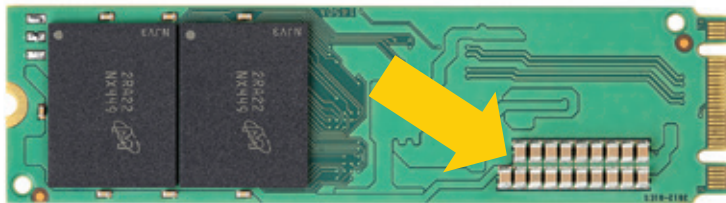


Figure 3: Ceramic Hold-Up Capacitors on Micron's M500 Client SSD

In SSDs used for client computing, data-at-rest protection is usually sufficient in the event of an unexpected power loss. Micron's M500 SSDs (and newer) accomplish data-at-rest protection using hardware controlled by Micron's proprietary firmware. Specifically, Micron SSD designers have added a special-purpose microcontroller, a holdup regulator, and a bank of standard ceramic capacitors (as shown in Figure 3) to detect and react to unexpected events at the power source. The microcontroller continually monitors the integrity of the source power (V_{CC}), and when V_{CC} dips below a certain level, the microcontroller signals the main SSD controller to stop new writes from the host, stop all new NAND operations (including background writes), and trigger the hold-up regulator to discharge the energy from the capacitors into the NAND components.

The hold-up capacitors provide only 1ms or so of hold-up time, which is enough time to place the signal in a lower-page program at its original level (recall Figure 1) prior to the start of the upper-page program. The capacitors are inexpensive and small enough to be deployed on SSDs with smaller form factor, such as Micron's mSATA and M.2 SSDs.

In addition to capacitive data protection, Micron's newest client SSDs employ a superior method of keeping track of data addressing information. Although Micron's SSDs use fairly large DRAM components, only a small amount of the DRAM is actually used to buffer user data. Rather, the DRAM

is used to manage the logical-to-physical address table (the FTL, as described earlier) in real-time. The new management method ensures that the FTL addressing is recorded and efficiently laid out such that if a power-loss event causes a loss of the address table, it can be rebuilt rapidly from information still stored in nonvolatile memory. Rebuilding the FTL table only adds a few seconds to the boot and can be complete in plenty of time

for the system BIOS or UEFI to “handshake” with the drive. The system may respond a bit more slowly than it would during a typical reboot, but is none-the-worse for wear after recovering from the unexpected power loss.

Countermeasures in Enterprise Computing

In enterprise computing, data-loss protection is considered to be much more critical than in client computing. Not only is it essential to protect data at rest, like in client SSDs, but it is also mandatory to protect data in flight. Any writes in progress must be completed and any data buffered in volatile memory must be committed to the NAND device and protected. As a result, Micron's enterprise SSDs implement a more significant

capacitive solution. As shown in Figure 4, instead of standard surface-mount ceramic capacitors, enterprise SSDs use larger tantalum capacitors, which store more energy to be discharged over a longer amount of time.

In Micron’s newer SSDs like the M500DC, the hold-up circuitry enables the SSDs to protect the stored lower-page programs, complete any in-progress writes to lower or upper pages, and dump any buffered writes from nonvolatile memory into the NAND. The hold-up circuitry also preserves enough time and energy to ensure that the FTL addressing table is properly saved to the NAND. This thorough

amount of data protection not only ensures data integrity in unexpected power-loss events, but it also enables the system designer to leave the SSD’s write cache enabled, giving a significant advantage in data throughput speeds.

Conclusion

Micron engineers are making every effort to protect data stored on Micron’s SSDs—even in the event of an unexpected power loss. This gives these devices significant robustness, even in the presence of uncertain power environments.

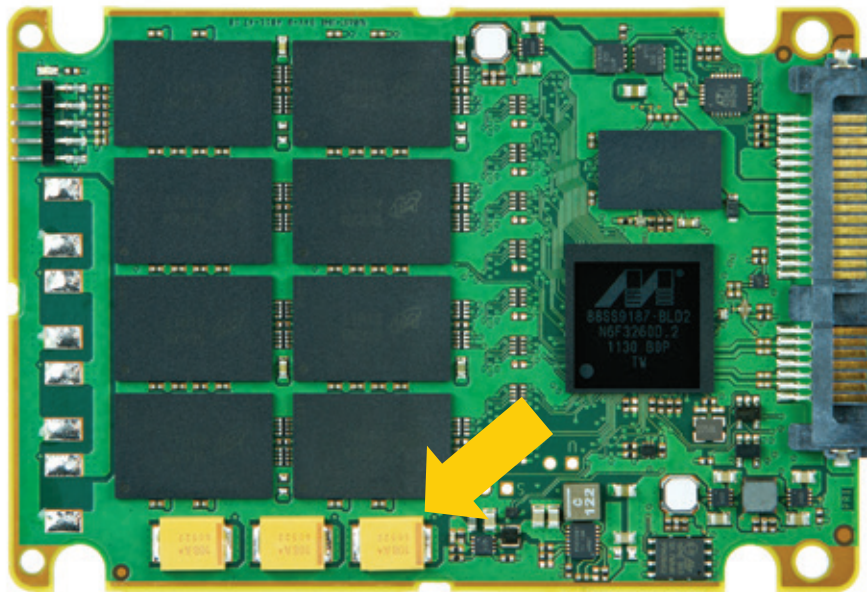


Figure 4: Example of Tantalum Hold-Up Capacitors on an Enterprise SSD

micron.com

Products are warranted only to meet Micron’s production data sheet specifications. Products and specifications are subject to change without notice. The information contained herein is provided on an “as is” basis without warranties of any kind.

Micron and the Micron logo are trademarks of Micron Technology, Inc. All other trademarks are the property of their respective owners. ©2014 Micron Technology, Inc. All rights reserved. 2/15