

Hardware Assets

Rugged Storage Products

Flash Disk Reliability Begins at the IC Level

At the flash-cell level, the die-level and the system level: to tackle the demands of rugged apps, rugged flash storage system designers must consider an army of factors affecting reliability.

Alan Fitzgerald, Chief Technology Officer
Adtron

A prominent device physicist and designer of flash memory once described using flash for storage as “the task of making order out of chaos”. Designers of the flash management system must turn the flash “chaos” into highly reliable storage systems rugged enough for demanding applications like the military. When used within an industrial duty, solid-state storage device tasked with emulating a hard disk or tape drive, there are several fundamental characteristics of flash memory that affect reliability.

Verification of flash quality is a shared responsibility between the flash chip manufacturer and flash storage system manufacturer. Once fully configured, the flash disk system must be thoroughly characterized by testing to exercise the flash failure modes to ensure meeting the flash disk reliability target. Figure 1 illustrates the process of characterizing flash.

Over the years flash die geometry, cell structure and addressing control structures have gone through revolutionary changes. As each new generation flash is produced, the process changeover causes lower yields and lower performance. Typically write endurance suffers the greatest decline in a technology transition. However, in most cases, the target endurance is reached and exceeded as the new fabrication process is fine-tuned. At each transition it is critical to characterize the

flash memory to ensure the flash matches the manufacturer’s specifications.

Die Sorting Issues

Testing and sort criteria characterize the quality of the flash memory at the chip factory to determine consumer or industrial quality. NAND and AND architectures—the most prominent flash flavor used for mass storage systems—are designed to allow bad blocks in a fully operational flash device. That is, there are more flash cells than required to meet the specified device capacity. Industry jargon of “mostly good” describes this technique.

At wafer test and die sort time, the assessment of good and best flash can depend in part on the percentage of mostly good flash. Across the area of flash die there is a distribution of “weak” and “strong” cells, and even more pronounced, across a wafer there will be a distribution of bad, weak and strong die. During wafer

probe it is possible to map the percentage of good cells and predict the quality of the flash device. It is not necessary for all the flash cells to be good. The bad cells are marked and the devices are sorted to consumer quality if the percentage good is moderate and to the industrial quality level if the percentage is close to 100%.

The sorting process will define the bad blocks so that the capacity of the die represents the available capacity and reflects the usable cells. It is the assumption of the sort that the devices with a low percentage of good cells will not be as reliable as the flash with the higher percentage of good cells. Table 1 shows each stage of the flash characterization process.

Defect Management

Decreasing cell geometry, change over from single to multi level cells and more dense control architectures, are making flash-based storage systems more

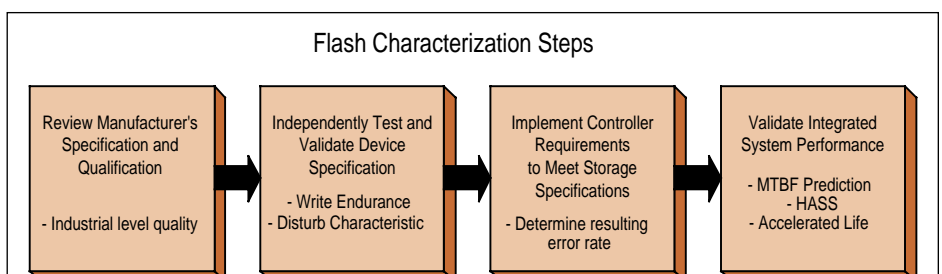


Figure 1

Depicted here is the process of characterizing flash. Once fully configured, the flash disk system must be thoroughly characterized by testing to exercise the flash failure modes to ensure meeting the flash disk reliability target.

Hardware Assets

dependent on defect management. With the confusion between consumer and industrial grade flash, it is more important than ever to ensure reliability through qualification, verification and production testing.

Only through testing can the flash quality be characterized and validated to ensure industrial quality and reliability. Low quality flash and poorly implemented defect management may not be apparent during simple functional test processes. Flash disk failures don't always reveal themselves during initial installation, and instead can show up after only a few months.

The flash disk manufacturer should qualify the flash and verify the effectiveness of the controller-based algorithms. These tests must include validation of each technology, write endurance, disturb and life tests. A reliability prediction method can give insight into the expected life of the flash disk from a component life perspective within a defined environ-

ment. The prediction method should be representative of the target environment, such as military through MIL-HDBK-217 or telecommunication through the Belcore prediction process.

Write Endurance

Early adopters of flash for storage suffered write endurance limits in many applications because the number of cycles was limited to about 10,000. Today it is common to find flash device write endurance over one million cycles and the use of defect management techniques to take care of infant failures. With modern write endurance and defect management techniques there are very few applications that are limited by write endurance.

The two most important management techniques to extend the write endurance life of flash are wear leveling and spare sectors. Spare sectors are just like spare tires, they are replacements for failed sectors. When a sector fails during a write or has a

correctable read error, it is retired and the data is written to the spare sector. This technique is also used in hard disk drives to overcome similar failures.

The inherent write endurance is effectively increased through the use of wear leveling techniques that even out the write cycles across a flash array used in a solid-state disk. The controller that manages the flash keeps track of the write usage within the flash array and moves the writing around the array to prevent concentrated write areas. Even static or read only areas are moved to accommodate the leveling of writes.

It's important to understand the inherent operation of the flash device with and without wear leveling and spare sectors. Characterizing the unaided write endurance of flash memory provides insight into the effectiveness of wear leveling and sector sparing. Qualification of any flash storage device should include this level of characterization—especially after every technology transition and significant fabrication process change.

Predicting Reliability and Accelerated Life Testing

There is no deterministic method to accurately predict mean time between failures (MTBF) of a component or system. Predicting performance within an application requires knowledge of the environment, the system characteristics and the acceptable failure rates. There are a number of factors that determine reliability influenced by not only the component used, but also the physical and electrical environment they're in.

Estimating component wear out within an application can help define the overall subsystem life within an application. Component failure rates can be calculated using standardized reliability models based on empirical testing and quality factors from several sources. MIL-HDBK-217, Belcore and British Telecom are all sources for reliability models and prediction methods. Each produce different results based on the empirical modeling for their defined environments.

One effective method for testing the predicted life of a component is accelerated life testing. The accelerated life testing process validates the predicted/expected life of a component or system based on a specific failure mode. The test process is based on the attribute of the storage system that is to be empirically

verified. This testing can include vibration, heat and cold, voltage, write or read frequency and duty cycle, data retention and others.

Product validation and general assessment of design robustness can be achieved through the use of HASS (Highly Accelerated Stress Screening). Typically these tests stress the unit under test with the application of vibration, heat, cold and temperature gradients. The test corners include high temperature and low temperature combined with vibration while operating to the last operating point and to the point of destruction. A failure analysis serves to illustrate a design weakness and recommendation for product improvement.

An important test is a long-term, operational test performed at the high end of the component specification. This test stresses the components, most significantly the integrated circuits, to accelerate their life. Silicon degrades with time and heat due to the effects of electro-migration, package seal leaks, oxide defects, ionic contamination and contact degradation. Using the Arrhenius equation with activation energy for each of these effects, an acceleration factor can be found and used to determine the accelerated life.

Disturb Errors

Disturb is one of those few technical adjectives that is very descriptive of this error event. Just like your neighbor playing the stereo too loudly, unintended changes to neighboring flash cells can occur during reading or writing. These activities can inject noise into adjacent cells that share control lines. With high duty cycle activity and adequate noise, internally and/or externally generated, the adjacent, non-participating flash cells can be modified by a write or read of another cell.

A disturb-induced error is discovered by the next read of the affected cells. The next read may be months away. To combat disturb errors, the defect management in the flash controller must be able to detect data errors and correct data errors on the fly. Error correction codes (ECC) are used to detect and correct this type of error. Since ECC algorithms are limited to correcting a finite number of bit errors, the flash quality must be adequate to keep bit errors within the correction limit so that disturb errors can always be corrected.

Some architectures provide higher immunity against disturb. The high-density NAND architecture flash achieves very effi-

Industrial Flash Qualification Process

Quality Process	Action	Purpose	Typical Value
Manufacturer specification and device qualification review	Analyze the process and results of the flash manufacturer's qualification process	Establish baseline for expected performance and required defect management algorithms	N/A
Write Endurance	Focused write/erase and distribution analysis, determine effectiveness of wear leveling and sector spare algorithms	Define the requirements for techniques such as wear leveling and sector sparing	Wide range from initial process of ~100K to mature processes >1 million cycles
Write and Read Disturb	Write and read throughout common control space to exercise sensitive areas	Define the requirements of error detection and correction	Corrected error rates are dependent on architecture < 10 ⁻¹⁴
Data Retention	Elevated temperature testing and Arrhenius equation prediction	Establish flash ability to retain/archive data without power	10 years to 20 years at 25 °C. Higher temperatures reduce time
MTBF Prediction	Statistical estimate of system MTBF based on empirical component tests within an environment	Establish reliability estimate of the system prior to empirical testing of the completed design	Dependent on the prediction environment MIL-HDBK-217: 250K hrs Belcore: >2 million hrs
HASS, Highly Accelerated Stress Screening	Application of temperature and vibration to determine last operating point and point of destruction	Design validation and assembly stress test to determine overall robustness	Typically looking for temperature corners 10 to 20 °C outside of specification while under vibration stress
Accelerated Life Test	Based on calculated Arrhenius acceleration factor, elevated temperature environment under continuous operation	Validate MTBF prediction and stress parts to find failure modes	Adtron has ongoing accelerated life testing with equivalent of 56 years of operation at 25 °C

Table 1

This table spells out the process, purpose and typical values at each stage of the flash characterization process.

cient use of silicon by sharing address control lines. This sharing makes disturb more prevalent in NAND than in other architectures, thus NAND flash controllers must incorporate stronger ECC algorithms to manage higher soft error rates.

The effects of disturb can be cumulative causing additional bit flips through high duty cycle disturb-inducing activity. Wear leveling can prevent the accumulation of bit flips by reading static data areas, thus improving the probability of detecting and correcting disturb errors, before too many bits are altered. Wear leveling is generally promoted as the method to improve write/erase endurance by spreading write activities across the flash memory array. However, by reading and moving static and low write activity blocks, bit flips are discovered, corrected with ECC and written to new locations.

Data Retention

The archive quality of flash data, i.e., the time that data can be stored without power, is the retention time. A typical value is 10 years at 25 °C. Testing processes vary to some degree, and should always be specified at the end of cell write endurance life to ensure worst case gate insulation characteristics. Industrial quality systems should have 10 years of data retention. Consumer requirements are significantly lower as it is likely that the pictures from last summer are already burned onto a CD.

Because of the long duration, 10 years or more, data retention life is not directly measured, but predicted using the Arrhenius equation at elevated temperatures and calculated using the acceleration factor. Typically for this test flash memory is stored at 150 °C, without power, and after the test period, data integrity is verified. From the high temperature testing, the Arrhenius curve will predict the period of time at room temperature for data retention.

Beyond Data Sheets

Managing the chaos of flash to produce a reliable and predictable storage system requires significant testing and validation of component behaviors. Simply accepting data sheet values and claims has prevented many designs from achieving their target reliability goal. Only through a

highly regimented program of characterization and validation is it possible to make the correct choice and subsequent implementation of a flash storage system.

Flash disks suitable for industrial and military use must be based on solid design techniques through the establishment of flash component quality and implementation of the necessary flash management techniques. This process must be repeated through every technology transition to ensure continued quality of the flash disk. No device specification should be assumed without a full qualification after every process change.

Of course the final validation is in the life testing and field life. Through accelerated testing (see Sidebar: "Predicting Reliability and Accelerated Life Testing") and analysis of field failure rates quality can be improved based on lessons learned. Flash disk manufacturers must have the corporate commitment to internal quality systems that support the continual improvement of quality. ■■



Adtron Corporation

Headquarters: Phoenix, AZ

European Offices: Switzerland and Denmark

Tel: U.S. 602-735-0300, Europe + 41-56-496-5640

Fax: U.S. 602-735-0359, Europe + 41-56-496-5648

marketing@adtron.com

www.adtron.com