



Nand Flash based Solid State Storage

Application Note / Glossary

BU: Flash Products
Revision: 1

File: Swissbit AppNote SSD Rev1.doc

Contents

1 SWISSBIT.....	3
2 INTRODUCTION.....	4
2.1 INTRODUCTION TO NAND FLASH MEMORY.....	4
2.1.1 Data Retention.....	6
2.1.2 Comparison to NOR Flash Memory.....	7
2.2 FLASH MEMORY ORGANIZATION.....	7
2.3 FLASH MEMORY CONTROLLER AND FIRMWARE.....	8
2.4 FLASH MEMORY MANAGEMENT.....	9
2.4.1 Blocks and Block Mapping.....	9
2.4.2 Error Code Correction (ECC).....	9
2.4.3 Bad Block Management.....	10
2.4.4 Wear Leveling.....	11
2.4.5 Endurance.....	11
2.4.6 Safe Power Loss Protection.....	12
3 C-300 SERIES COMPACT FLASH.....	13
3.1 POWER DOWN.....	13
4 SOLID STATE DRIVES (SSD) VERSUS HARD DISK DRIVES (HDD).....	14
4.1 FEATURE POSSIBILITIES AND EMERGING APPLICATIONS.....	14
4.2 REQUIREMENTS ARE CHANGING.....	14
5 GLOSSARY.....	15
5.1 NAND FLASH.....	15
5.2 FIRMWARE / CONTROLLER FEATURES.....	16
5.3 PRODUCTS / INTERFACES.....	16
5.4 TRANSFER MODES.....	17
5.5 FILE SYSTEM.....	17
6 REFERENCES.....	17

1 Swissbit

Embedded and Industrial systems require a variety of memory and storage solutions. In contrast to the desktop and portable computer and consumer market, the products are exposed to critical environmental conditions, are designed for longer life cycles, and therefore must be highly reliable. Swissbit's industrial Solid State Drive (SSD) family covers all relevant interfaces, including CompactFlash (CF), Secure Digital (SD), IDE / Parallel ATA (PATA), Serial ATA (SATA) and Universal Serial Bus (USB), in combination with state of the art Flash handling utilizing the qualities of single level cell (SLC) NAND Flash.

Product Reliability

Product development according to stringent design rules, critical component control, and extensive product qualification procedures, ensure overall electrical and mechanical robustness of Swissbit's Flash product line. All products are offered in commercial (0°C to +70°C) and industrial (-40°C to +85°C) temperature ranges. Only highly reliable SLC NAND Flash is used and our Flash handling provides features such as built-in error correction, bad block management, static and dynamic wear leveling, and power loss protection. For our latest product revisions we offer product lifetime calculations for special use cases. We enable our customers to access bad block counts and the status of write/erase cycles in order to monitor wear leveling. This enables the calculation of the remaining product lifetime in the field.

Partners and Competences

Since its inception, Swissbit has many years of experience in the memory business. We have long term relationships with major semiconductor manufacturers and work closely together with our controller vendors.

Technical Support

With our engineering expertise we provide support during product qualification and can help end customers to integrate our products into their applications. We can also help design products from the ground-up or adjust existing products to meet special requirements. Firmware or software support, as well as customer system failure analysis, are well within our capabilities.

2 Introduction

Flash memory has significantly replaced various historic storage media. As examples, some 5 years back people were using film in non-digital cameras, computers were offering floppy disk drives, portable CD players played your favorite tunes, and mobile phones were incapable of taking pictures, browsing the web or synchronizing with a PC. In recent years, consumers have become accustomed to using a large variety of NAND Flash memory devices for data storage applications. Users of mobile phones, digital cameras, and other gadgets are familiar with the terms SD card, USB-stick, CF cards and the like. Such kinds of small and handy memory systems have become everyday commodities among consumers.

Requirements for such consumer applications are normally determined by

- Cost
- Compatibility to a plethora of host devices
- Fast transfer rates for reading and writing of data
- Storage capacity
- Conformity to standards
- Ease of use

Memory cards or USB sticks that fulfill these requirements can meanwhile be found in almost any supermarket or grocery store around the world. As the above features are not too difficult to obtain, it has become a challenge to list the variety of vendors who manufacture or sell such memory systems.

Have you ever experienced the sharp sound of your hard disk when it crashes? If so, you may specifically like the features of a solid state disk (SSD). Flash based solutions offer many advantages such as being faster, more power efficient, more rugged than rotating media, and being more easily integrated together with other chips in system design and production flows. Tremendous technology advances have decreased manufacturing costs and reduced prices significantly. These advances however are more and more demanding for very intelligent control functions.

Requirements for highly reliable Flash memory storage systems, especially those not removable, such as SSD's, embedded Flash, or industrial products, are more demanding than removable consumer grade cards:

- Good mechanism to detect and correct errors inherent to Flash memories
- Efficient algorithms to maximize lifetime
- Tools to predict lifetime or monitor systems' status and health
- Ability to implement customer-specific features

2.1 Introduction to NAND Flash Memory

In many implementations, memory cells are based on CMOS floating-gate transistors. Each cell can represent one or more bits by reading out one or multiple levels of its electrical charge at the word line. This charge is changed using the Fowler-Nordheim tunneling or tunnel effect. Electrons are removed from or trapped in the floating gate. By applying a positive potential to either the bit line or the word line, the charge is changed and the cell is either erased or programmed respectively.

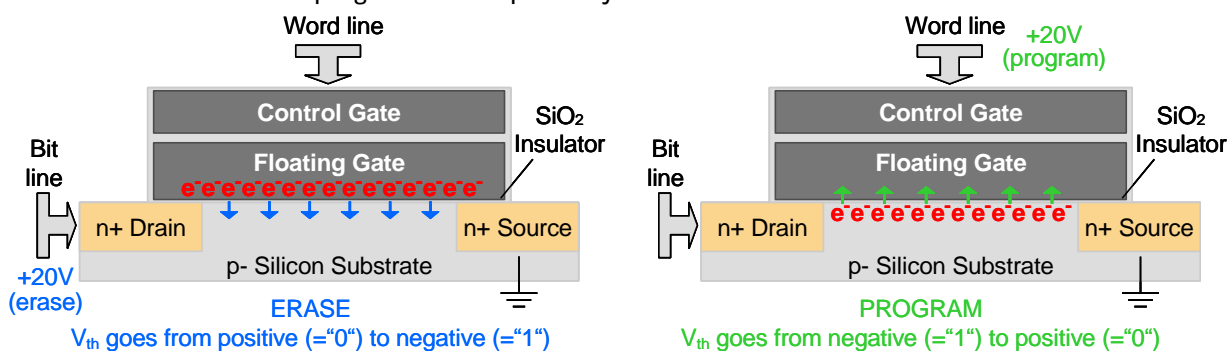


Figure 1: Program & Erase for an Example of a Floating Gate Memory Cell

So called Single-Level Cells (SLC) store or represent one bit. Charges (or threshold voltages V_{th}) between about $-1V$ and $-3V$ could be considered as representing the value "1" and charges between $+1V$ and $+3V$ as value "0".

So called Multi-Level Cells (MLC) store or represent two or more bits of information per cell. While the erase process and the erased level for Multi Level Cells is similar to that of SLC, for the programming or programmed state different charge levels of the floating gate have to be achieved. This is performed by applying Incremental Step Pulse Programming (ISPP). Again a positive potential in the range of 15 to 25 V is applied to a word line, but in short pulses. After each pulse a check is made to determine whether a desired charge level is reached. As soon as the desired charge has been reached and confirmed, the cell is considered programmed and represents two (or more) bits of logical information. Therefore, the writing process of MLC takes a little longer.

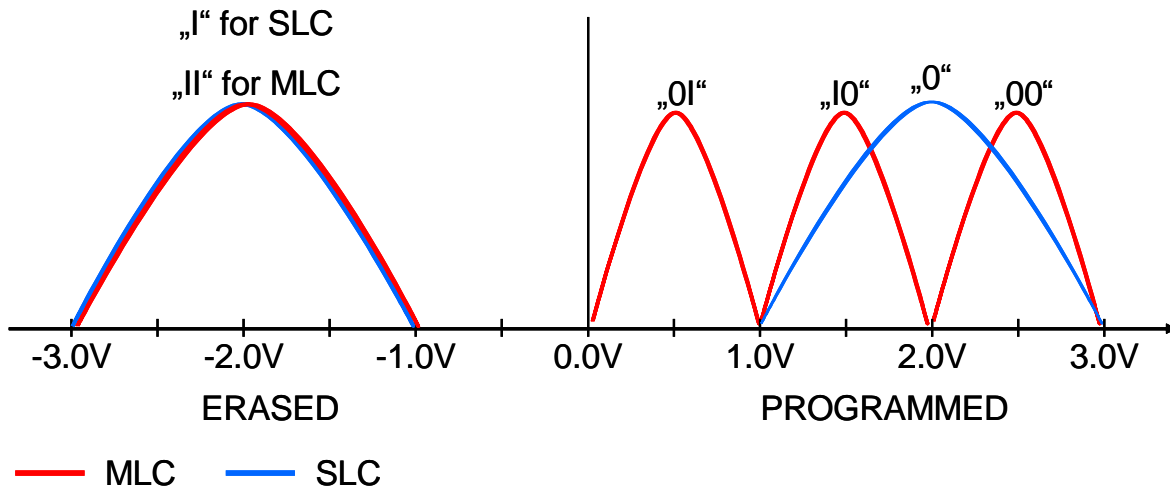


Figure 2: Example of possible voltage levels and related logical bit value of SLC and MLC compared

Writing and erasing imposes stress on the individual cell so that cell ages or wear down over program/erase cycles. In effect, the charge levels change over time. While programming time decreases, erasing takes an increasing amount of time. Eventually charge levels exceed the defined thresholds and errors occur. As thresholds for MLC must be closer together, production variances have a greater effect and errors occur more readily and/or earlier over time.





Figure 3: Example of possible voltage levels for SLC and MLC changing over time

Additionally, charge levels might change due to external conditions such as extreme heat or magnetism. While the cell itself is not damaged permanently, the bit value might have changed and a read error might occur.

Some more recent Flashes have the capability of recognizing the systematic change in behavior or in the voltage level so that not the difference to a starting reference voltage, but the inability to differentiate the relative difference to other voltage levels produces the read errors.

Understanding technical limitations such as those resulting from process shrinks, it is also important to know that effects such as capacitive-coupling of internal signals and floating-gate-to-floating-gate noise can cause the cell's threshold voltage to shift by programming neighboring cells. This is one of the reasons that many believe there to be a physical barrier to shrinking geometries in the 30 nm region or slightly below.

2.1.1 Data Retention

While a cell's "quality" over time is mainly dependent on the number of write/erase cycles, the quality of a cell's voltage level or the data retention might also deteriorate over time. Often Flash data sheets specify data

retention in the area of 10 years, while real-life data retention might be subject to all kinds of intrinsic factors such as heat, magnetism, traffic on neighboring cells and more. On the positive side, remapping or wear leveling might improve the data retention, as data is actually moving across physical blocks and the last time a bit is physically written differs from the last time it was logically written.

2.1.2 Comparison to NOR Flash Memory

Cells in NAND Flash are arranged in arrays of between 8 and 32 cells. Unlike in NOR Flash, the individual cells are not connected to the bit line. For this reason, NOR Flash requires more area and is slower to program and erase, but on the positive side, NOR Flash achieves better random access times and can be programmed by byte.




Figure 4: NAND Flash array of 16 cells

2.2 Flash Memory Organization

When looking at how a Flash memory is organized the smallest logical/administrative unit is a sector. Each sector contains 512 bytes plus an overhead area (traditionally 16 bytes). One to eight sectors are then grouped into pages, in the range of 512 to 4,096 bytes per page. At the next level of hierarchy blocks can include 32 pages of 512 bytes for example, or more recently 64 pages of 2,048 bytes. Blocks therefore also have a defined number of sectors, currently between 16 and 512 and there are 1 to 8 thousand blocks per chip.

Physical Block Addresses																	
Block 0					Block 1					Block n							
Page n			Page 1		Page 0		Page n			Page 1		Page 0		Page n			
Sector 0	Sector 1	Sector n	Sector 0	Sector 1	Sector n	Sector 0	Sector 1	Sector n	Sector 0	Sector 1	Sector n	Sector 0	Sector 1	Sector n	Sector 0	Sector 1	Sector n

Figure 5: Organization of a Flash Memory

Writing or programming is done at the page level. Also, programming requires pages to be pre-erased. Once pre-erased, programming or writing can be performed one page at a time or by addressing a sector within a page that is pre-erased or "empty". Depending on Flash memory type, pages can be accessed up to 4 times, when writing a sector.

Erasing is performed at the block level. As blocks are the "management" or "administrative" units, blocks will wear out as memory cells break down after a number of erase cycles. When defective, these blocks will be considered "bad blocks".

Flash memories from different manufacturers vary greatly with respect to bus width, number of blocks, block size, page size, die count, and programming capabilities, including caches. The following figure shows the organization of a typical NAND Flash device available today.

Example Memory organization:

Sector Size	512	Byte
Sectors/Page	8	
Pages/Block	64	
Page Size	4	KByte
Block Size	256	KByte
Blocks / Die	4096	
Dies/ Chip	2	
Flash Capacity	2	GByte

Figure 6: Example Organization of a Flash Memory (16 Gbit, dual die, SLC, 4K pages)

All of these characteristics pose some interesting challenges to the controller and its firmware:

- How to re-write to areas already containing data?
- How to maximize product lifetime with only a limited amount of erase cycles available?
- How to ensure data transfer integrity?

2.3 Flash Memory Controller and Firmware

The complex nature of Flash cells and their organization demands reliable, high performance control functionality. Flash Controllers of all kinds consist of an interface to the Flash memory, a processor and a host interface.

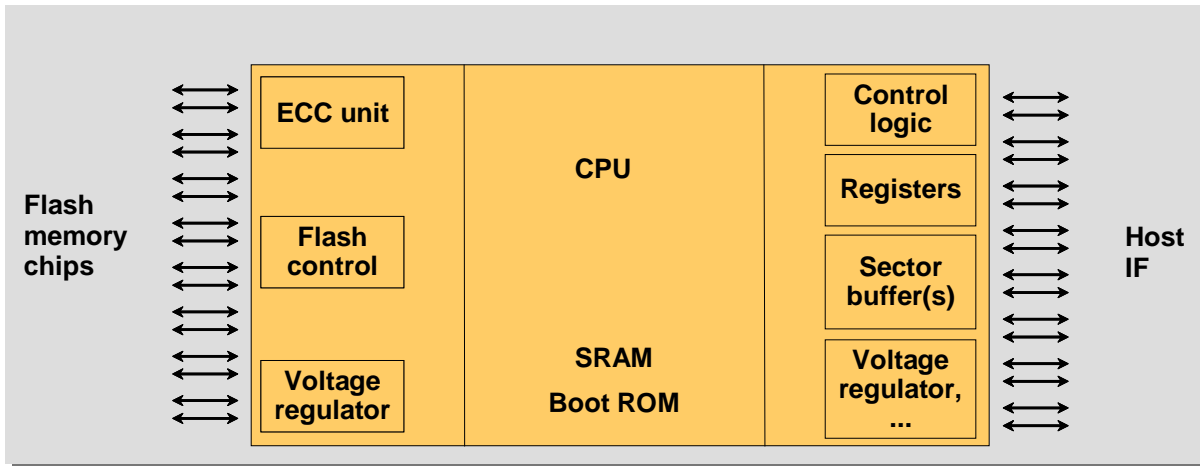


Figure 7: Generic Block Diagram Solid State Memory Solution

The controllers are based on a CPU together with dedicated hardware blocks, including an error correcting code (ECC) unit, buffers, Flash and host interface control logic.

The firmware, stored in the Flash memory, is flash device, target application, and host interface specific. All tasks with respect to Flash, data management, and data transfers between Flash and host are implemented either in hardware or in software. The Flash controller boot-up is accomplished by using firmware that is stored within the Flash memory of the product. Other solutions might store firmware in the ROM of the controller. Therefore, manufacturers are able to provide different products or feature sets all based on identical product hardware. Also, because the firmware is copied into the Flash in a so-called 'pre-formatting' process after the storage product has been assembled, firmware can be updated in the field or immediately before delivery.

2.4 Flash Memory Management

Several algorithms and concepts are used to address the questions initially posed in re-writing to areas already containing data, maximizing Flash life time, and ensuring data transfer integrity.

2.4.1 Blocks and Block Mapping

Logical block addresses (LBA) including the related static sector address or information are mapped correspondingly to physical block addresses (PBA). A table is maintained by the controller or firmware, translating requests for the particular LBA to its corresponding PBA. Logical blocks are distributed across all available Flash chips and 'good' blocks, e.g. logical block 0 might correspond to a physical block at chip 1 block 2, and logical block 1 might correspond to a physical block on chip 3, block 3.




Figure 8: Logical to Physical Block Addressing

2.4.2 Error Code Correction (ECC)

The ECC is generally responsible for ensuring the quality of data being read, by adding information during the writing phase to restore partially corrupted data. ECC algorithms encode a larger data array based on certain assumptions or criteria in order to detect and repair errors. To get into the logic, one could start with a simple error detection method such as a parity bit, check sum or cyclic redundancy check (CRC). A check sum is as simple as adding all bits and comparing results. This is an easy and efficient method in terms of data overhead for error detection but doesn't help to repair data. Therefore more complex encoding is needed. So called syndromes contain more information that not only detect but can also help to restore data. Encoding additional information requires more overhead and calculation power. The Hamming code, a linear error-correcting code based on a parity matrix, has been widely used in telecommunications.

Data Package	Data Bit A	Data Bit B	Data Bit C	Data Bit D	Parity Bit 1 C1=A+B+C	Parity Bit 2 C2=A+B+D	Parity Bit 3 C3=A+C+D
P1 sent	1	0	0	1	1	0	0
P2 sent	0	1	1	0	0	1	1

Figure 9: Example of a Hamming Code encoding

Data Package	Data Bit A	Data Bit B	Data Bit C	Data Bit D	C1' = A+B+C	C2' = A+B+D	C3' = A+C+D	C1=C1'? C2=C2'? C3=C3'
P1 received	0	0	0	1	0	1	1	ERROR
P2 received	0	1	1	0	0	1	1	O.K.

Figure 10: Example of a Hamming Code used for Error detection

Identified False Bit	C1' = A+B+C	C2' = A+B+D	C3' = A+C+D
----------------------	-------------	-------------	-------------

None	True	True	True
Data Bit A	False	False	False
Data Bit B	False	False	True
Data Bit C	False	True	False
Data Bit D	True	False	False
Parity Bit C1'	False	True	True
Parity Bit C2'	True	False	True
Parity Bit C3'	True	True	False

Figure 11: Table of Equations and results for all possible error bits

By performing the following calculations and comparing the results to the above table, the false bit can be identified and corrected as necessary:

$$C1' = A \text{ XOR } B \text{ XOR } C = 1 \rightarrow \text{False since the value of } C1' \text{ received is } 0$$

$$C2' = A \text{ XOR } B \text{ XOR } D = 0 \rightarrow \text{(False since the value of } C2' \text{ received is } 1)$$

$$C3' = A \text{ XOR } C \text{ XOR } D = 0 \rightarrow \text{(False since the value of } C3' \text{ received is } 1)$$

As all values are false, Data Bit A was received wrong and can now be corrected. This example shows how an ECC algorithm can work and also is an example as to the overhead required in terms of additional data that needs to be stored and transmitted.

As for the Hamming Code, 3 bits had to be added to 4 bits of information. This ratio might not be considered very efficient. Reed–Solomon code works similarly for multi-bit symbols. A (255,235) Reed–Solomon code RS (255,235) – specifies a total block length of 255 bytes (or symbols); 235 bytes used for information and 20 check bytes. The check bytes are calculated in a similar manner to the 3 check bits in the Hamming code example above and are appended to the end of the data block. Reed–Solomon codes are much more complex however, and require a significant amount of arithmetic and logical processing. Reed–Solomon codes are very code rate efficient. Within Flash handling, error correction is done on a bit/byte level within sectors, i.e. a 4–Bit Reed–Solomon ECC, refers to the capability of correcting 4 bytes, that is a maximum of 32 Bit, within a sector of 512 (+16) data byte.

Therefore, Flashes specify a so-called spare area for each page, for example when we talked about a 512 byte or 4Kbyte page before, we actually refer to e.g. 528 and 4224 Bytes per page where the additional memory is not available to users but used by an ECC to store overhead information. This also results in the fact that the ECC can be limited by the spare area of each page provided on the Flash.

Both Reed–Solomon and the Hamming Code are good for burst errors as are more common for data transmission and telecoms. Since Flash cell errors occur more sporadically other algorithms are appearing in products, such as a more flexible binary BCH¹, for which parity and syndrome calculation as well as correction can all be done in hardware. For this concept, six bits are correctable for a 16 byte overhead area per 512 data bytes, or 14 bits correctable for a 32 byte (30 byte or larger) overhead area. In general, since ECC algorithms can either be implemented in software or in hardware there is a trade-off between flexibility to support the many different Flash memory cell architectures and performance of the ECC. Because of the very small cell structures in new MLC Flash, ECC tasks are of increasing importance and also increasingly specific to cells' architectures.

The Swissbit S–2x0, C–300, P–100, and X–100 Series of Flash products all use a 4 symbol Reed–Solomon ECC. Parity and syndrome generation are performed on-the-fly (in hardware) during the data transfer to/from the Flash. Correction itself is performed in software, with Galois field arithmetic support through special CPU instructions. Additionally and in order to make sure the ECC has worked correctly, a cyclic redundancy check (CRC) is added in order to check if the resulting “corrected data” is indeed correct. Correctable errors are corrected without further notice. Uncorrectable errors in user data are notified to the host setting an “UNC bit” in the error register.

2.4.3 Bad Block Management

In order to ensure that all blocks used for data storage are functional, but knowing that blocks wear out or become defective for several reasons, bad block management is necessary to retire all questionable blocks. During initial pre-formatting, Flash memories are tested and bad blocks originally marked by the Flash manufacturers are subsequently mapped out. Bad blocks are entered into a bad block table referring to physical block addresses. A pool of spare blocks is defined and used for dynamic bad block replacement, where blocks from the pool are used to replace blocks that produce errors when erased or programmed. This is called ‘bad block re-mapping’. Considerations as to when blocks should be considered “bad” range from recognizing write or erase fails, reported by the flash itself, to algorithms based on bit errors recognized by the ECC during

reading. During the Flash's life within an application and while managed by an external controller, the bad block table is maintained and extended by any block going bad as it is mapped out or 'retired'. As soon as all spare blocks are consumed, can set into read only mode.

The probability of bit errors, or the bit error rate, increases with the number of write/erase cycles, with smaller process geometries, and also for MLC compared to SLC architectures. Therefore, a more powerful ECC (e.g. other than specified in the Flash data sheet) can extend the lifetime of blocks over the specified number of write/erase cycles. Currently, the ECC is implemented as part of the controller and is responsible for the quality of reading from Flash. The quality of writing and erasing is determined by circuits within a Flash. This is important, as both *do not* necessarily communicate. Therefore, it is possible that the Flash reports a write/erase failure too early, when indeed the external ECC would be capable of correcting the errors. On the other hand, the controller on the Flash does not consider information regarding the bit error rate of reading. Finally, bit errors when reading do not necessarily mean a block is indeed bad, simply that the information is corrupted to some extent. Re-writing the information into another block might cure the system altogether if in fact the data corruption was due a temporary external effect and not to a worn out block. Erasing that involves writing the value "1" for SLC or "11" for MLC is of great significance in managing the Flash. It is the most strenuous of activities for the gate hence the one where most errors occur. At the same time it is the easiest to verify since the result for each cell is known and can be easily checked, such as building a check sum for instance. Therefore, pre-erasing can function almost like a gate keeper, weeding out bad blocks, if the process is managed in the right way.

2.4.4 Wear Leveling

The algorithms do balancing the use of all blocks, thus guarantying maximum lifetime of products. Furthermore, algorithms relating to and integrated in the wear leveling can help to avoid capacitive or parasitic-coupling, floating-gate-to-floating-gate noise, electron contamination and such by optimizing the spread of cells being addressed.

The wear leveling algorithms are erase count triggered. All blocks are classified into so-called 'wear level classes'. Whenever a block is erased, a counter is increased. If the erase counter reaches a defined threshold, the block's wear level class is increased. By comparing wear level classes of blocks to be used, the load is balanced throughout all blocks on a Flash memory chip. By trying to ensure that as few blocks as possible reach their end-of-life before others, and together with defining an adequate size of the spare block pool, the products life can be maximized and data integrity be ensured.

We provide vendor specific commands, providing, amongst other features, information regarding wear-level classes of blocks. Using the command "read wear level count", system developers can read block usage statistics. This information combined with defined thresholds can serve as an endurance status indicator.

2.4.5 Endurance

Floating gates "wear out" and after a certain amount of program/erase cycles, a gate becomes defective. A certain amount of defective gates can be compensated by the ECC. As soon as the amount of defective gates (bit error rate) within a block is larger than the ECC can handle, a block is mapped out as a bad block.

Usually, endurance of SLC Flash is specified with 100K program/erase cycles, while endurance of MLC Flashes reach up to 10K program/erase cycles. Typically a reference to some kind of ECC is specified along with it, which could be 1-Bit for SLC and 4-Bit for MLC.

The following calculation only holds true under the assumption that perfect wear leveling is applied and all static data that has been programmed is still involved in the pool of blocks. As a very first approximation, the total "cycle potential" could be estimated by multiplying the number of blocks by the specified maximum program/erase cycles. For example, with 8K blocks and 100K program/erase cycles, we could assume a "cycle potential" of a single SLC Flash to be 800 million [block write/erases].

In actuality, estimating a products' lifetime is more complicated because one would need to know how this "cycle potential" is "consumed". Let's look at a scenario where we would re-write a file equal to a block size of 256 kB. With 8.000 blocks we could write 100.000 times to each block that is 800 million write erase cycles in total. Physically that file would have been located 100K times in each block.

On the other hand, looking at a scenario of rewriting a file that is equal in size to the total available memory, in our example that is 2 GByte, wear leveling or page wise programming would not have any effect. This file could be written 100K times before reaching the specified limit.

In order to translate a “cycle potential” into a lifetime in years it must be known how frequently data is written. Again, taking a very theoretical example of a block size file and assuming programming happens every second, the 800 million cycles would result in a life expectancy of about 25 years (800 million / $1 \times 60 \times 60 \times 24 \times 365$). For MLC this would still be 3 years of continuously writing that file each second. The administrative overhead in terms of wear leveling etc. for Swissbit products is only in the range of 1 –2%. The following tables can give a rough estimate of the order of magnitude regarding this endurance approximation.

File Size equal to:	e.g.	2 GB, SLC Flash	2 GB, MLC Flash	
One Block	256 kB	800,000,000	80,000,000	# cycles
8 Blocks	2 MB	100,000,000	10,000,000	# cycles
80 Blocks	20 MB	10,000,000	1,000,000	# cycles
800 Blocks	204.8 MB	1,000,000	100,000	# cycles
Capacity	2,048 MB	100,000	10,000	# cycles

Figure 12: Example of an endurance approximation in number of write/erase cycles

File Size equal to:	e.g.	write frequency	2 GB, SLC Flash	2 GB, MLC Flash	
One Block	256 kB	1 x each second	25	3	Years
8 Blocks	2.048 MB	1 x each minute	190	19	Years
80 Blocks	20.48 MB	1 x each minute	19	2	Years
800 Blocks	204.8 MB	1 x each hour	114	11	Years
Capacity	2,048 MB	1 x each hour	11	1	Years

Figure 13: Example of an endurance approximation in years based on assumed data write frequencies

Without wear leveling, endurance estimation heavily depends on what address might be written to. Even a very small file, if written to the same address over and over again, might destroy blocks fast and in the extreme worst case the overall endurance is very close to the specified write/erase cycles.

Another important factor to endurance is the quality of the ECC compared to the statistical occurrence of read errors. Usually, a kind of ECC is specified along with the number of write/erase cycles e.g. 1-Bit for SLC and 4-Bit for MLC. Using a better ECC can increase these numbers significantly.

2.4.6 Safe Power Loss Protection

Flash memory is often used in removable storage applications or battery operated devices where a robust and reliable power source cannot be guaranteed. A user may remove the memory at any time and under these conditions security of data is of paramount importance. Swissbit uses a concept in order to ensure data integrity when transferring or writing data. By using certain buffer blocks, information is written in a way that minimizes the delta between an old and a new state. The data system is coherent at all times.

Upon a sudden power fail, the controller is reset and the Flash is immediately write-protected. A log of the most recent Flash transactions is kept, where entries are made just before any programming to the Flash. Should the last entry of the log be corrupted, the controller recovers the last valid entry. This minimizes data loss due to power failures and data corruption at the physical layer is prevented completely. Should power loss happen at the very same time when data is written to the Flash, this data might get lost. In no case, however, will the overall data system be corrupted. Swissbit performs extensive power cycling tests to all products verifying no data corruption due to power failure.

3 C-300 Series Compact Flash

As an example of a more complex product, the C-300 Series offers two channels, “Enhanced Direct Flash Access” each with a 4 symbol Error Correcting Code (ECC) unit capable of correcting 4 Bytes in a 512 bytes sector. Up to 16 Flash memory chips (eight chip-selects per channel) can be directly connected. The 32-Bit RISC microprocessor can be scaled, running at clock frequencies from 10 MHz to 60 MHz using a trim able internal oscillator. The CPU offers 16 Kbytes internal Boot ROM, 20 Kbytes internal SRAM, four 512 Byte sector buffers and 256 Byte PCMCIA attribute memory. Performance with SLC Flash memories above 40 MBytes/s reading and over 30 MBytes/s writing is achievable. With automatic power down mode during wait periods for host data or Flash memory operation completion and automatic sleep mode during host inactivity periods, a current of less than 200 µA can be achieved.

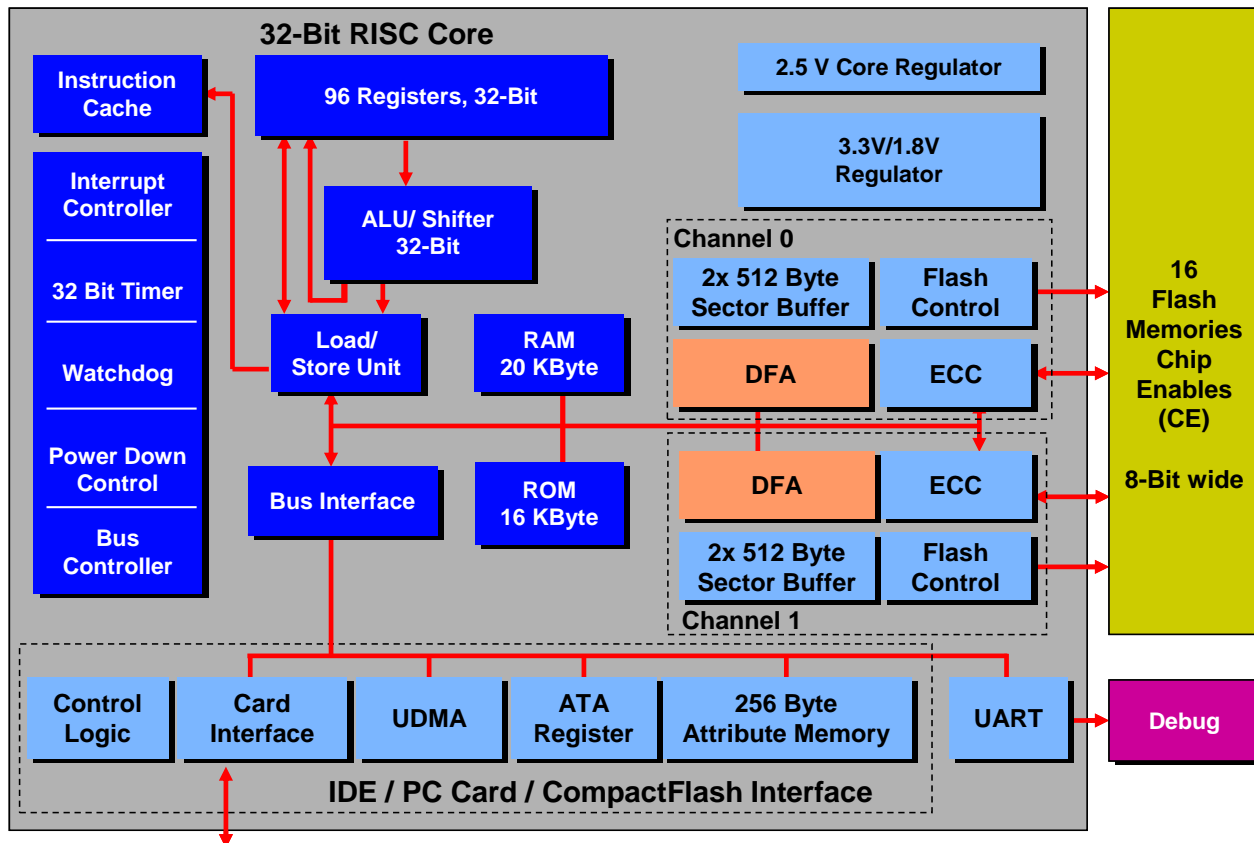


Figure 14: Block Diagram of C-300 Series used Memory Controller for CF card (Hyperstone F3 SSD controller)

The C-300 Series is specifically designed for applications that require high-speed CF cards for the industrial market. It supports automatic sensing of PCMCIA or True-IDE host interface mode, and is compliant with PCMCIA 2.1, PC Card ATA, CF 3.0, memory mapped or I/O operations, fast ATA host-to-buffer transfer rates, PIO mode 6, MDMA mode 4, and UDMA mode 4 in true-IDE mode.

Power down

- Sector mapping table in Flash and RAM
- On every update this table is stored to the Flash
- Worst case 16 sectors of data could be lost, on power on old data before the write would be restored and file system integrity would be ensured
- Static data cannot be destroyed, even if data is written into the same block during power off / failure
- Data cannot be destroyed, if a block is copied to a new position during power off / failure
- FW /cards are tested in power cycle tests over days where power down situations are simulated at random.

4 Solid State Drives (SSD) versus Hard Disk Drives (HDD)

Following Flash cards and removable data storage systems, the latest killer application for Flash is considered to be solid state drives (SSD), a substitute for hard disk drives or fixed storage devices. Flash has some significant advantages over hard disk drives most of which due to the absence of moving parts, power consumption and robustness to physical impact:

- Performance: Latency times basically due to the inherent mechanical operations of HDD are not necessary for SSD and seek times or access times are about 30 times faster. Sustained writing using interleaving can be scaled to some extent by both HDD and SSD. However, again due to the mechanical issues, this is more difficult for HDDs. Finally the HDD has a slight advantage at sustained reading. Performance in real life scenarios also depends on the level of fragmentation of the data when using HDDs. Have you used your defrag tool lately? Or have you ever performed this for a 500GB drive? With SSDs, defragmentation is an inherent feature of the logical to physical translation and wear-leveling.
- Power Consumption:
 - Power Idle: HDD 0.8–5.0 W; SSD 0.035–3.0 W
 - Power Read/Write: HDD 5.0–10.0 W; SSD: 0.325–5.0 W
- Reliability: Again due to the mechanics HDD most significant disadvantage is its reliability and sensitivity due to shock
- Scalability of Form Factors: Solid State Storage can be implemented in many different form factors from microSD or eMMC to 3.5" drives.
- Choosing the optimal Host Interfaces: Solid State Storage can be based on several host interfaces including but not limited to SD, MMC, USB, S-ATA, P-ATA/IDE, CF, while HDDs are mainly driven by the PC market using mainly S-ATA today.
- Endurance (Write/Erase Cycles)
 - HDD: NA
 - SSD: 100,000 → 5 million

4.1 Feature Possibilities and Emerging Applications

Several features can be realized aside of standard features. Cards can be defined as read-only, content can be protected, and writing to a card can be limited or defined in number, e.g. write-once, -twice, x times. Event triggered write-protection is possible, as is making cards function only in certain devices. Proprietary hardware format can make use of standards such as SD/MMC, ATA, or IDE together with firmware comprising a disk-on-board. Requirements for qualification, tight quality control, or simply longer life cycles compared to consumer products might easily be addressed, and result in specific products or in-sourcing.

4.2 Requirements are changing

New generations of Flash will continue to improve performance, increase capacity, or drive cost per MByte down. These advances include technology shrinks (60, 50, 40 nm), larger page sizes (2KB, 4KB), several Flash interfaces (8-Bit, 16-Bit, ONFI, HL-NAND), and finally cell technologies with different write or erase cycle performance and maximum life time (SLC, 2-Bit MLC, 4-Bit MLC). Furthermore, new packaging and manufacturing processes will continue to drive overall system cost down and capacity up as well. All of these advances have an impact on the controller or firmware implementation and as a result, controllers and firmware features will become more important as an enabling and competitive factor of applications.

While ensuring highest performance, reliability, enabling highest capacity, guaranteeing data integrity and maximum life-time are most important features determined by controllers, cost of controllers together with firmware represent only a small fraction of the overall systems' cost.

5 Glossary

5.1 NAND FLASH

Type of Flash typically used for data storage applications as opposed to NOR Flash used in embedded systems, typically for program storage. Gets its name from the NAND Cell structure used.

- PBA (Physical Block Address)
 - Actual physical address of Flash block the sector resides in.
- BLOCK
 - Inside the Flash the largest unit of data storage. Typically 128K or 256K for large block Flash. Currently an 8Gbit SLC Flash has 8192 Blocks
- PAGE
 - Inside the Flash device, equal to 4 (2K Page) or 8 sectors (4K page)
- Page Overhead Area
 - Each Flash page has extra bytes (currently 16) per 512 byte sector. This is typically used by the controller for overhead items such as CRC, parity, logical block number etc. This brings the physical sector size to 528 bytes.
- Flash Internal Copy Page (Copy-back)
 - A mechanism where data from one Flash block is copied to another without leaving the Flash and without Controller intervention.
- Two Plane Flash and Commands
 - The newer Flash has its data blocks divided into odd and even planes. There are two plane commands with such Flash, that allows simultaneous operations on one block from each plane. So there is a two-plane program, erase etc. Using these commands allows significant increase in speed. Currently most Flash has this feature.
- Flash Cache Program
 - This allows writing of new data while the previous data is being programmed into the Flash array.
- SLC (Single Level Cell) FLASH – Expensive
 - 1 bit per Cell.
 - Reliable with 100K typical Program/Erase Cycles.
 - Data Retention is 10 years.
 - Good for Solid State Disk (SSD) applications.
 - An ECC of 1 bit (in 512 Bytes) is generally recommended.
- MLC (Multi-Level Cell) FLASH – Cheap
 - 2 or more bits per cell; allows larger density in the same package.
 - 10K program/Erase Cycles before EOL.
 - Good for commodity applications
 - 4 bits (in 512 Bytes) or more of ECC
 - As the process shrinks below 50 nm, the expectations are that 12 bit ECC will be required.
- SDP, DDP, QDP Flash
 - Single Die package or Mono Die, Dual Die Package, Quad Die Package
 - SDP and DDP packages typically have 1 Chip Enable
 - QDP packages typically have 2 chip enables
- NAND Flash Manufacturers
 - Samsung, Hynix, Micron, Intel, Nymonix, Toshiba, etc.

5.2 Firmware / Controller Features

- ECC (Error Correcting Code)
 - NAND Flash requires the use of an ECC. Parity bits are added so that data errors can be detected and corrected, within the limits of the ECC used.
- CRC (Cyclical Redundancy Check)
 - Used for the Ultra DMA protocol to check the validity of the data. Also used to backup the ECC.
- Wear Leveling
 - Since NAND Flash has a limited life, its important that all parts of the Flash experience similar usage. Wear Leveling assures that.
- Bad Block Remap
 - The process of replacing a Bad Block with a good one from the spare block pool.
- LBA (logical block address)
 - This term defines the addressing of the device as being by the linear logical mapping of sectors
- DFA (Direct FLASH Access)
 - A means of data transfer between the controller and the Flash without intervention from the controller.
- Interleave
 - A method that uses the busy time from a Flash to device(s) to simultaneously send program/erase commands to another Flash device(s)
- Interleave Factor
 - The number of Flash devices that are simultaneously serviced for Program/Erases. For the F3 this is 1,2 or 4. Interleaving significantly increases write speeds.
- Anchor Block
 - Chip o Block o. All permanent information stored here as well as operating firmware. It is the "Anchor" point for administrative information. The Anchor Sector is Sector 0 of the Anchor Block.

5.3 Products / Interfaces

- Universal Serial Bus (USB)
 - Universal Serial Bus (USB) is a serial bus standard to interface devices to a host computer. USB was designed to allow many peripherals to be connected using a single standardized interface socket and to improve the plug-and-play capabilities by allowing hot swapping, that is, by allowing devices to be connected and disconnected without rebooting the computer or turning off the device. Other convenient features include providing power to low-consumption devices without the need for an external power supply and allowing many devices to be used without requiring manufacturer specific, individual device drivers to be installed.
<http://en.wikipedia.org/wiki/USB>
- USB Flash Drive (UFD)
 - A USB flash drive consists of a NAND-type flash memory data storage device integrated with a USB (universal serial bus) interface. USB flash drives are typically removable and rewritable, much shorter than a floppy disk (1 to 4 inches or 2.5 to 10 cm), and weigh less than 2 ounces (56 g). Storage capacities typically range from 16 MB to 64 GB with steady improvements in size and price per gigabyte. Some allow more then 100 thousand write or erase cycles and have 10-year data retention, connected by USB 1.1 or USB 2.0.
http://en.wikipedia.org/wiki/Usb_flash_drive
- COMPACT FLASH (CF)
 - 50 Pin Parallel Interface Card, a subset of PCMCIA. Can operate in PC Card modes or TRUE IDE mode (like a Disk Drive). Typically used in professional digital Cameras
- ATA / PATA / IDE
 - AT Attachment Standard – ATA defines the physical, electrical, transport, and command protocols for the internal attachment of storage devices. A parallel Disk Drive Interface standard often referred to as PATA and IDE.
- Serial ATA (SATA)
 - Serial version of ATA. Capable of data transfer rates > 200 MB/S
- Secure Digital (SD)
 - Serial Interface memory card with Content Protection available with CPRM keys. Typically 4 bits, extended to 8 bits. Royalties.
- Multi-Media (MMC)
 - Similar to SD without any royalties. 8 bit capability

5.4 Transfer Modes

- SPI Mode
 - The SPI mode consists of a secondary communication protocol that is offered by Flash-based SD Memory Cards. This mode is a subset of the SD Memory Card protocol, designed to communicate with a SPI channel, commonly found in Motorola's (and lately a few other vendors') microcontrollers.
- PIO modes
 - Parallel Input/Output. Basic ATA/IDE data transfer mode using I/O methods. Slower than DMA modes. PIO0 – PIO6 are defined.
- DMA (Direct Memory Access)
 - A means of data transfer between device and host memory without processor intervention
 - Multi-Word (DMA) Modes: Basic DMA transfer modes for IDE drives. Allows transfers of multiple data words without processor intervention. Often abbreviated MDMA.
 - Ultra DMA (UDMA) Modes: Faster mode of DMA using double edge clocking. UDMA5. Speeds of 100 MB/S are possible in UDMA5 mode, 133MB/S for UDMA6 mode. CRC is included in the data stream.

5.5 File system

- SECTOR
 - Basic unit of data transfer = 512 bytes and smallest unit of data inside the Flash device. Physical sector size is actually 528 bytes to allow overhead information.
- FAT
 - File Allocation Table or FAT is a computer file system architecture originally developed by Bill Gates and Marc McDonald in 1976/1977. It is the primary file system for various operating systems including DR-DOS, OpenDOS, freeDOS, MS-DOS, OS/2(v1.1), and Microsoft Windows (up to Windows Me). For floppy disks (FAT12 and FAT16 without long filename support) it has been standardized as ECMA-107 and ISO/IEC 9293.
- FAT32
 - In order to overcome the volume size limit of FAT16, while still allowing DOS real mode code to handle the format without unnecessarily reducing the available conventional memory, Microsoft implemented a newer generation of FAT, known as FAT32, with cluster values held in a 32-bit field, of which 28 bits are used to hold the cluster number, for a maximum of approximately 268 million (2²⁸) clusters. This allows for drive sizes of up to 8 terabytes with 32KB clusters, but the boot sector uses a 32-bit field for the sector count, limiting volume size to 2 TB on a hard disk with 512 byte sectors.
- NTFS
 - NTFS (New Technology File System) Is the standard file system of Windows NT, including its later versions Windows 2000, Windows XP, Windows Server 2003, Windows Server 2008, and Windows Vista.
 - NTFS supersedes the FAT file system as the preferred file system for Microsoft's "Windows"-branded operating systems. NTFS has several improvements over FAT and HPFS (High Performance File System) such as improved support for metadata and the use of advanced data structures to improve performance, reliability, and disk space utilization, plus additional extensions such as security access control lists (ACL) and file system journaling.

6 References

www.swissbit.com
www.hyperstone.com
<http://en.wikipedia.org/>

© Swissbit AG, certain parts authorized by Hyperstone GmbH