

Rino Micheloni
Luca Crippa
Alessia Marelli

Inside NAND Flash Memories

 Springer

Inside NAND Flash Memories

Rino Micheloni • Luca Crippa • Alessia Marelli

Inside NAND Flash Memories

 Springer

Rino Micheloni
Integrated Device Technology
Agrate Brianza
Italy
rino.micheloni@ieee.org

Luca Crippa
Forward Insights
North York
Canada
luca.crippa@ieee.org

Alessia Marelli
Integrated Device Technology
Agrate Brianza
Italy
alessiamarelli@gmail.com

ISBN 978-90-481-9430-8 e-ISBN 978-90-481-9431-5
DOI 10.1007/978-90-481-9431-5
Springer Dordrecht Heidelberg London New York

Library of Congress Control Number: 2010931597

© Springer Science+Business Media B.V. 2010

No part of this work may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, microfilming, recording or otherwise, without written permission from the Publisher, with the exception of any material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work.

Cover design: eStudio Calamar S.L.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)




Fig. 1.10. NAND device characteristics (Source: Forward Insights)




Fig. 1.11. Applications for multi-bit per cell NAND Flash memories (Source: Forward Insights)




Fig. 1.16. Solid state drive versus hard disk drive (Source: Toshiba Corp., Forward Insights)

Figure 1.16 details the major components of an SSD and HDD. The HDD, being based on storage in a spinning magnetic platter, requires an actuator and actuator arm to move the head to the appropriate sector to be read or written. This movement of the read head results in extremely long latencies in the milliseconds and a DRAM buffer is used to hide the seek time. The HDD controller, particularly if it is a system-on-chip, incorporates the processor, servo control logic, interface, error correction code, disk sequencer and buffer controller.

In contrast, the SSD contains no mechanical parts and consists of a few major components: NAND Flash memory, SSD controller, connector, DRAM, PCB and passives. In addition, because of the small size of NAND Flash memory, the SSD form factor is not limited to standard 1.8, 2.5 or 3.5" HDD form factors but can also come in module form.

NAND-based SSDs started out in industrial and military where ruggedness and reliability are a priority. In 2006, the first SSD for personal computers was introduced followed by SSDs for enterprise computing.

Table 1.1 summarizes the NAND Flash/SSD usage models in the different market segments. There are several ways NAND Flash or SSDs may be utilized. Firstly, it can act as an alternate storage device, in many cases replacing a HDD for latency or reliability improvement and power consumption reduction. An optimized OS or file-system is required to obtain the full benefits of the SSD.

Secondly, the SSD may act as an I/O accelerator. I/O acceleration is primarily used in the enterprise computing environment. A software driver is required for this approach.

Thirdly, the SSD/NAND Flash may be part of the tiered system memory and replace part of the DRAM. Due to the better economics of NAND Flash versus DRAM, this allows an increase in system memory storage capacity with the

benefits of reduced power consumption. To manage this tiered system memory, the kernel memory manager needs to be designed to manage the Flash memory.

A comparison of SSD and HDD in PCs (Fig. 1.17) shows that SSDs are superior to HDDs in all aspects except for cost and capacity. The cost and capacity are related since the major impediment to higher SSD capacities is the high cost per gigabyte.

Table 1.1. NAND usage model

NAND Usage Model	Consumer	Enterprise	Industrial	Military
Alternate Storage	√	√	√	√
I/O Accelerator		√		
Extended System Memory	√	√		

In the enterprise space, high throughput and low access latency are critical. Enterprise HDDs employ several techniques to improve system throughput and decrease the access latency. High performance 15 k rpm HDDs may be organized in a RAID configuration with the data being striped across multiple disks thereby increasing the bandwidth of the total system. Adding HDDs increase system bandwidth but result in increased space, power and cooling requirements.

“Short-stroking” the HDDs whereby data is placed on the outer tracks of the disk to reduce the seek time of the mechanical heads reduces system access latencies. However, this entails that only part of the HDD is being used to achieve the required performance resulting in excess storage capacity than would otherwise be necessary. As illustrated in Fig. 1.18, one SSD can replace several striped, short-stroked HDDs.




Fig. 1.17. 2.5" SSDs versus HDDs in PCs – performance characteristics (Source: Forward Insights)

The most popular Flash memory cell is based on the *Floating Gate* (FG) technology, whose cross section is shown in Fig. 2.1. A MOS transistor is built with two overlapping gates rather than a single one: the first one is completely surrounded by oxide, while the second one is contacted to form the gate terminal. The isolated gate constitutes an excellent “trap” for electrons, which guarantees charge retention for years. The operations performed to inject and remove electrons from the isolated gate are called program and erase, respectively. These operations modify the threshold voltage V_{TH} of the memory cell, which is a special type of MOS transistor. Applying a fixed voltage to cell’s terminals, it is then possible to discriminate two storage levels: when the gate voltage is higher than the cell’s V_{TH} , the cell is on (“1”), otherwise it is off (“0”).

It is worth mentioning that, due to floating gate scalability reasons, charge trap memories are gaining more and more attention and they are described in Chap. 5, together with their 3D evolution.




Fig. 2.1. (a) Floating gate memory cell and (b) its schematic symbol

2.2 NAND memory

2.2.1 Array

The memory cells are packed to form a matrix in order to optimize silicon area occupation. Depending on how the cells are organized in the matrix, it is possible to distinguish between NAND and NOR Flash memories. This book is about NAND memories as they are the most widespread in the storage systems. NOR architecture is described in great details in [1].

In the NAND string, the cells are connected in series, in groups of 32 or 64, as shown in Fig. 2.2. Two selection transistors are placed at the edges of the string, to ensure the connections to the source line (through M_{SL}) and to the bitline (through M_{DL}). Each NAND string shares the bitline contact with another string. Control gates are connected through wordlines (WLs).




Fig. 2.2. NAND string (left) and NAND array (right)

Logical pages are made up by cells belonging to the same wordline. The number of pages per wordline is related to the storage capabilities of the memory cell. Depending on the number of storage levels, Flash memories are referred to in different ways: SLC memories store 1 bit per cell, MLC memories (Chap. 10) store 2 bits per cell, 8LC memories (Chap. 16) store 3 bits per cell and 16LC memories (Chap. 16) store 4 bits per cell.

If we consider the SLC case with interleaved architecture (Chap. 8), even and odd cells form two different pages. For example, a SLC device with 4 kB page has a wordline of 65,536 cells.

Of course, in the MLC case there are four pages as each cell stores one *Least Significant Bit* (LSB) and one *Most Significant Bit* (MSB). Therefore, we have:

- MSB and LSB pages on even bitlines
- MSB and LSB pages on odd bitlines

where V_{THN} indicates the threshold voltage value of the NMOS M_N .

At this point, the M_N and M_P transistors are switched off. C_{BL} is free to discharge. After a time T_{VAL} , the gate of M_N is biased at $V_2 < V_I$, usually 1.6–1.4 V.

If a T_{VAL} time has elapsed long enough to discharge the bitline voltage under the value:

$$V_{BL} < V_2 - V_{THN} \quad (2.2)$$

M_N turns on and the voltage of node OUT (V_{OUT}) becomes equal to the one of the bitline. Finally, the analog voltage V_{OUT} is converted into a digital format by using simple latches.

A more detailed analysis of read techniques is presented in Chap. 8.

Program

Programming of NAND memories exploits the quantum-effect of electron tunneling in the presence of a strong electric field (Fowler–Nordheim tunneling [2]). In particular, depending on the polarity of the electric field applied, program or erase take place. Please refer to Chap. 3 for more details.

During programming, the number of electrons crossing the oxide is a function of the electric field: in fact, the greater such field is, the greater the injection probability is. Thus, in order to improve the program performances, it is essential to have high electric fields available and therefore high voltages (Chaps. 11 and 12). This requirement is one of the main drawbacks of this program method, since the oxide degradation is impacted by these voltages.

The main advantage is the current required, which is definitely low, in the range of nanoamperes per cell. This is what makes the Fowler–Nordheim mechanism suitable for a parallel programming of many cells as required by NAND page sizes.

The algorithm used to program the cells of a NAND memory is a *Program & Verify* algorithm (Chaps. 3 and 12): verify is used to check whether the cell has reached the target distribution or not.

In order to trigger the injection of electrons into the floating gate, the following voltages are applied, as shown in Fig. 2.5:

- V_{DD} on the gate of the drain selector
- $V_{PASS,P}$ (8–10 V) on the unselected gates
- V_{PGM} (20–25 V) on the selected gate (to be programmed)
- GND on the gate of the source selector
- GND on the bitlines to be programmed
- V_{DD} on other bitlines

The so-called *self-boosting* mechanism (Chap. 3) prevents the cells sharing the same gate with the programmed one from undergoing an undesired program. The basic idea is to exploit the high potentials involved during programming through

the parasitic capacitors of the cell, in order to increase the potential of the region underneath the tunnel oxide (inset of Fig. 2.5).




Fig. 2.5. NAND string biasing during program and parasitic capacitors contributing to the self-boosting inhibit

When the bitlines are driven to V_{DD} , drain transistors are diode-connected and the corresponding bitlines are floating. When $V_{PASS,P}$ is applied to the unselected wordlines, parasitic capacitors boost the potential of the channel, reducing the voltage drop across the tunnel oxide and, hence, inhibiting the tunneling phenomena.

As the memory cells are organized in a matrix, all the cells along the wordline are biased at the same voltage even if they are not intended to be programmed, i.e. they are “disturbed”. Two important typologies of disturbs are related to the program operation: the *Pass disturb* and the *Program disturb*, as shown in Fig. 2.5: their impact on reliability is described in Chap. 4.

Erase

NAND memory is placed in a triple-well structure, as shown in Fig. 2.6a. Usually, each plane has its own triple-well. The source terminal is shared by all the blocks: in this way the matrix is more compact and the multiplicity of the structures which bias the iP-well is drastically reduced.

The electrical erase is achieved by biasing the iP-well with a high voltage and keeping grounded the wordlines of the sector to be erased (Fig. 2.6c). Therefore, NAND technologies don’t need negative voltages. Again, the physical mechanism is the Fowler–Nordheim tunneling. As the iP-well is common to all the blocks,

erase of unselected blocks is prevented leaving their wordlines floating. In this way, when the iP-well is charged, the potential of the floating wordlines raises thanks to the capacitive coupling between the control gates and the iP-well (Fowler–Nordheim tunneling inhibited).

Figure 2.6b sketches the erase algorithm phases. NAND specifications are quite aggressive in terms of erase time. Therefore, Flash vendors try to erase the block content in few erase steps (possibly one). As a consequence, a very high electric field is applied to the matrix during the *Electrical Erase* phase. As a matter of fact, erased distribution is deeply shifted towards negative V_{TH} values. In order to minimize floating gate coupling (Chap. 12), a *Program After Erase* (PAE) phase is introduced, with the goal of placing the distribution near the 0 V limit (of course, guaranteeing the appropriate read margin).

Typical erase time of a SLC block is about 1–1.5 ms; electrical erase pulse lasts 700–800 μ s.

Technology shrink will ask for more sophisticated erase algorithms, especially for 3–4 bit/cell devices. In fact, reliability margins are usually shrinking with the technology node: a more precise, and therefore time consuming, PAE will be introduced, in order to contain the erased distribution width. In summary, erase time is going to increase to 4–5 ms in the new NAND generations.

Chapter 12 contains a detailed description of the whole erase algorithm.




Fig. 2.6. (a) NAND matrix in triple-well; (b) erase algorithm; (c) erase biasing on the selected block

2.2.3 Logic organization

NAND memory contains information organized in a specified way. Looking at Fig. 2.7, a memory is divided in pages and blocks. A block is the smallest erasable unit. Generally, there are a power of two blocks within any device. Each block contains multiple pages. The number of pages within a block is typically a multiple of 16 (e.g. 64, 128). A page is the smallest addressable unit for reading and writing. Each page is composed of main area and spare area (Fig. 2.8). Main area can range from 4 to 8 kB or even 16 kB. Spare area can be used for ECC (Chap. 14) and system pointers (Chap. 17) and it is in the order of a couple of hundreds bytes every 4 kB of main area.

Every time we want to execute an operation on a NAND device, we must issue the address where we want to act. The address is divided in row address and column address (Fig. 2.9). Row address identifies the addressed page, while column address is used to identify the bytes inside the page. When both row and column addresses are required, column address is given first, 8 bits per address cycle. The first cycle contains the least significant bits. Row and column addresses cannot share the same address cycle.

The row address identifies the block and the page involved in the operation. Page address occupies the least significant bits.




Fig. 2.7. NAND memory logic organization



Fig. 2.8. NAND page structure

All these issues cause a severe limitation to the maximum capacity of the card; in addition external components, like voltage regulators and quartz, cannot be used. In other words, the memory controller of the card has to implement all the required functions.

The assembly stress for small form factors is quite high and, therefore, system testing is at the end of the production. Hence, production cost is higher (Chap. 15).




Fig. 2.32. Block diagram of a SSD




Fig. 2.33. Functional representation of a Flash card (or SSD)

For a more detailed description of Flash cards, please, refer to Chap. 17. SSDs are described in Chap. 18.

Figure 2.33 shows a functional representation of a memory card or SSD: two types of components can be identified: the memory controller and the Flash memory components. Actual implementation may vary, but the functions described in the next sections are always present.

2.4.1 Memory controller

The aim of the memory controller is twofold:

1. To provide the most suitable interface and protocol towards both the host and the Flash memories
2. To efficiently handle data, maximizing transfer speed, data integrity and information retention

In order to carry out such tasks, an application specific device is designed, embedding a standard processor – usually 8–16 bits – together with dedicated hardware to handle timing-critical tasks.

For the sake of discussion, the memory controller can be divided into four parts, which are implemented either in hardware or in firmware. Proceeding from the host to the Flash, the first part is the host interface, which implements the required industry-standard protocol (MMC, SD, CF, etc.), thus ensuring both logical and electrical interoperability between Flash cards and hosts. This block is a mix of hardware – buffers, drivers, etc. – and firmware – command decoding performed by the embedded processor – which decodes the command sequence invoked by the host and handles the data flow to/from the Flash memories.

The second part is the Flash File System (FFS) [6]: that is, the file system which enables the use of Flash cards, SSDs and USB sticks like magnetic disks. For instance, sequential memory access on a multitude of sub-sectors which constitute a file is organized by linked lists (stored on the Flash card itself) which are used by the host to build the File Allocation Table (FAT).

The FFS is usually implemented in form of firmware inside the controller, each sub-layer performing a specific function. The main functions are: *Wear leveling Management*, *Garbage Collection* and *Bad Block Management*. For all these functions, tables are widely used in order to map sectors and pages from logical to physical (Flash Translation Layer or FTL) [7, 8], as shown in Fig. 2.34.

The upper block row is the logical view of the memory, while the lower row is the physical one. From the host perspective, data are transparently written and overwritten inside a given logical sector: due to Flash limitations, overwrite on the same page is not possible, therefore a new page (sector) must be allocated in the physical block and the previous one is marked as invalid. It is clear that, at some point in time, the current physical block becomes full and therefore a second one (Buffer) is assigned to the same logical block.

The required translation tables are always stored on the memory card itself, thus reducing the overall card capacity.




Fig. 2.34. Logical to physical block management

Wear leveling

Usually, not all the information stored within the same memory location change with the same frequency: some data are often updated while others remain always the same for a very long time – in the extreme case, for the whole life of the device. It's clear that the blocks containing frequently-updated information are stressed with a large number of write/erase cycles, while the blocks containing information updated very rarely are much less stressed.

In order to mitigate disturbs, it is important to keep the aging of each page/block as minimum and as uniform as possible: that is, the number of both read and program cycles applied to each page must be monitored. Furthermore, the maximum number of allowed program/erase cycles for a block (i.e. its endurance) should be considered: in case SLC NAND memories are used, this number is in the order of 100 k cycles, which is reduced to 10 k when MLC NAND memories are used.

Wear Leveling techniques rely on the concept of logical to physical translation: that is, each time the host application requires updates to the same (logical) sector, the memory controller dynamically maps the sector onto a different (physical) sector, keeping track of the mapping either in a specific table or with pointers. The out-of-date copy of the sector is tagged as both invalid and eligible for erase. In this way, all the physical sectors are evenly used, thus keeping the aging under a reasonable value.

Two kinds of approaches are possible: Dynamic Wear Leveling is normally used to follow up a user's request of update for a sector; Static Wear Leveling can also be implemented, where every sector, even the least modified, is eligible for re-mapping as soon as its aging deviates from the average value.

Garbage collection

Both wear leveling techniques rely on the availability of free sectors that can be filled up with the updates: as soon as the number of free sectors falls below a given threshold, sectors are “compacted” and multiple, obsolete copies are deleted.

This operation is performed by the Garbage Collection module, which selects the blocks containing the invalid sectors, copies the latest valid copy into free sectors and erases such blocks (Fig. 2.35).

In order to minimize the impact on performance, garbage collection can be performed in background. The equilibrium generated by the wear leveling distributes wear out stress over the array rather than on single hot spots. Hence, the bigger the memory density, the lower the wear out per cell is.




Fig. 2.35. Garbage collection

Bad block management

No matter how smart the Wear Leveling algorithm is, an intrinsic limitation of NAND Flash memories is represented by the presence of so-called *Bad Blocks* (BB), i.e. blocks which contain one or more locations whose reliability is not guaranteed.

The *Bad Block Management* (BBM) module creates and maintains a map of bad blocks, as shown in Fig. 2.36: this map is created during factory initialization of the memory card, thus containing the list of the bad blocks already present during the factory testing of the NAND Flash memory modules. Then it is updated during device lifetime whenever a block becomes bad.




Fig. 2.36. Bad Block Management (BBM)

Toggle mode adds DQS data strobe signals; RE is used to trigger the read cycle as done in asynchronous interface; DQS is used to strobe the data on both edges.

7.1.5 High density systems: density, power and performance

Today, SSDs are offered in densities up to 64–128 GB and higher. To satisfy Enterprise requirements of endurance and quality, producers are forced to use SLC devices or MLC with trailing-edge technology. As of today, a 32 GB SSD can be built using eight packages (e.g. TSOP), each containing four 16 Gbit-sub 50 nm generation SLC NAND. The system can be built using four independent data channels. This means that two TSOP packages are connected to a single channel with a total of eight dies hanging on the same data bus which is loaded roughly by 50 pF for each data line.

Some considerations are fundamental when designing a system:

- Which is the minimum needed configuration in terms of channels number and which is the operating frequency to achieve the required throughput?
- How much power is dissipated in the system? Is it sustainable?

Due to the channel capacitance, each time a single NAND die is being written or read, the entire capacitance (that can be as high as 50 pF in the example above) of the data lines is driven.

I/O power consumption in a DDR system can be derived as [26]:

$$P = 9 \cdot \eta \cdot f \cdot C \cdot V^2 \quad (7.1)$$

where η is the bit activity ratio, f is the DDR frequency, C is the capacitance of a single line and V is the supply voltage of the interface.

For the system described $\eta = 0.5$, $C = 50$ pF and $V = 3.3$ V. We can derive the power consumption as a direct function of throughput (Fig. 7.8).




Fig. 7.8. I/O power as a function of channel throughput

Of course, the latch I1–I2 can also be used without exploiting the three-state option. In this case, M_{EN} , M_{SA} and the latch must be correctly sized so that the latch unbalancing always takes place when V_{SO} goes below $V_{DD} - V_{THP}$.




Fig. 8.28. ABL sense amplifier and its timing diagram

8.4 ABL versus interleaving architecture

In this section a comparison between ABL and the interleaving architecture is presented [25].

Such sensing circuit can also be used to implement the read operation for other types of coding, like the one shown in Fig. 10.15. As described in the table of Fig. 10.16 (ENX column indicates the enabling signal), the read of the upper-page is done by two SROs: the former using V_{READ1} and the latter using V_{READ2} . Read of the lower-page needs three SROs, at V_{READ1} , V_{READ2} and V_{READ3} . In this case as well, latch L1 is busy doing the read operation while latch L2 performs the data-out, which means that even this structure belongs to the cache read category.

The implementation of the read operation for the type of coding shown in Fig. 10.7 can be easily derived: since the coding of the upper-page and the coding of the lower-page are swapped, it is sufficient to swap the labels between the upper-page and the lower-page in the table of Fig. 10.16.




Fig. 10.16. MLC with different distribution coding and cache read

It should be noted that using the solution shown in Fig. 10.13 it is possible to read both pages using only three SROs, while the example shown in Fig. 10.16 requires five SROs.

It should always be possible to discriminate the position of the threshold voltage using only three SROs. For instance, the solution shown in Fig. 10.17 implements the simultaneous read of the two pages: the two latches, L1 and L2, work in parallel during SROs. At the end of three SROs, the lower-page is on node A, and the upper page is on node OUT.




Fig. 10.17. MLC page buffer for reading two pages

10.2.3 MLC program/verify operations

During program operation, the same program pulse (at a gate voltage equal to V_{PGM}) is applied to all the cells belonging to the same wordline. It is responsibility of the sensing circuit to verify and, if necessary, inhibit those cells which have reached the desired distribution. In particular, the cell is inhibited and it does not change its threshold V_{THR} any longer if during the program step the corresponding bitline is biased to a voltage which is high enough to switch off selection transistor M_{BLS} (Fig. 10.6). The voltage applied to the bitline is typically equal to either V_{DD} or $V_{DD} - V_{THN}$. On the other hand, the cell increments its threshold voltage if during the program step the bitline is biased to a voltage which is low enough to


switch on M_{BSL} so that it can act as a pass transistor. Typically, the bitline is biased to ground and the situation just described is shown in Fig. 10.18.

In order to switch-off the transistor M_{BSL} , its gate is biased to the lowest possible voltage, considering that M_{BSL} has to be on when the bitline is forced to ground: typical values lie between 1 and 2 V.

Program of the upper-page is the most complex operation and its analysis allows us to understand the fundamental elements which must be present inside the sensing circuit. First of all, in all the coding cases previously seen, two information are needed in order to program the upper-page properly:

- Which is the initial distribution: the state of the lower-page must be known and stored in the sensing circuit. If the lower-page has already been programmed, the information must be recovered from the cell.
- Which is the target distribution: this information, given by the value of the bit of the upper-page, is provided by the user and it must be stored in the sensing circuit.

Since two pieces of information must be stored, it is evident that at least two latches are needed inside the sensing circuit. Furthermore, the circuit has to be interfaced to a logic-unit capable of loading the information that program needs (the data-load described in Sect. 10.3).



- $PB(n)$ discharge $BL(n)$ to ground: cell $M_{i,n}$ programmed
- $PB(n+1)$ charge $BL(n+1)$ to V_{DD} : cell $M_{i,n+1}$ inhibited

Fig. 10.18. Cells program and inhibit

14 Error correction codes

T. Zhang¹, A. Marelli² and R. Micheloni³

14.1 Introduction

As we have seen in Chap. 13, redundancy covers manufacturing defects while ECC takes care of the failures during the life of the device. This chapter deals with error correction codes applied to NAND Flash memories. In fact, when the memory is placed in its final application, different reasons for errors (see Chap. 4) can damage the written information so that it could happen that the read message is not equal to the original anymore [1].

Before proceeding on, it is necessary to introduce some basic definitions.

The error probability p , also known as raw bit error rate (BER), is defined as:

$$p = \frac{\text{Number of bit errors}}{\text{Total number of bits}} \quad (14.1)$$

The Chip Error Probability (CEP) is defined as:

$$CEP(p) = \frac{\text{Number of chip errors}(p)}{\text{Total number of chips}} \quad (14.2)$$

Considering a memory chip formed by B blocks of A bits, the CEP before ECC (CEP_{in}) is given by:

$$CEP_{in}(p) = 1 - (1 - p)^{AB} \quad (14.3)$$

Assuming to use a block correction code of t errors, the CEP_{out} is

$$CEP_{out}(p) = 1 - (1 - b)^B \quad (14.4)$$

where b is the failure probability of the block to which the ECC is applied

$$b = 1 - [P_0 + P_1 + P_2 + \dots + P_t] \quad (14.5)$$

¹ Rensselaer Polytechnic Institute, tzhang@ecse.rpi.edu

² Integrated Device Technology, alessiamarelli@gmail.com

³ Integrated Device Technology, rino.micheloni@ieee.org

P_i is the probability of i errors in a block of A bits

$$P_i = \binom{A}{i} p^i (1-p)^{A-i} \quad (14.6)$$

In conclusion we have that

$$CEP_{out}(p) = 1 - \left[(1-p)^A + \binom{A}{1} p(1-p)^{A-1} + \dots + \binom{A}{t} p^t (1-p)^{A-t} \right]^B \quad (14.7)$$

Figure 14.1 shows the graphs of CEP_{in} (indicated as “No ECC”) and CEP_{out} as a function of p for a system with 4 Gbit memory, 2 kB block and ECC able to correct one, two, three, four or five errors. Assuming a raw BER of 10^{-9} , 1-error ECC gives a Chip Error Probability of 100 ppm, while a 2-errors ECC recovers five orders of magnitude.




Fig. 14.1. Graph of the Chip Error Probability for a system with 4 Gbit memory. The CEP parameters for correction codes of one, two, three, four and five errors are given

14.2 Mathematical background

Error correction codes add redundant terms to the message, such that, on reception, it is possible to detect the errors and to recover the message that has most probably been transmitted.

Correction codes are divided into two well separated groups: block codes and convolutional codes. The first ones are so defined because they treat the information as separated blocks of fixed length, in contrast with convolutional codes which work on a continuous flow of information.

Another important distinction is between hard decision codes and soft decision codes. This distinction is not based on the structure of the code itself but on the way the information is treated by the code. A binary hard decision code treats all the data in a digital way, i.e. “0” or “1”; in other words, the analog information is converted into digital format using one fixed reference level. On the other hand, a soft decision code uses also reliability information about the decision, e. g. a “0” is read with a 90% reliability or a “1” is read with a 10% reliability.

14.2.1 Block codes

Schematically, a block m of k symbols is encoded in a block c of n symbols ($n > k$) and written in a memory. Inside the memory, different sources may generate errors e (see Chap. 4), so that the block message r is read. The block r is then decoded in d by using the maximum likelihood decoding strategy, so that d is the message that has most probably been written [2–4].

A *Code C* (Fig. 14.2) is the set of codewords obtained by associating the q^k messages of length k of the space A to q^k words of length n of the space B in an univocal way. A code is defined as *linear* if, given two codewords, also their sum is a codeword. When a code is linear, encoding and decoding can be described with matrix operations.

Definition 14.1 G is called *generator matrix of a code C* when all the codewords are obtainable as a combination of the rows of G .

Each code has more than one generator matrix, that is all its linear combinations. Each code has infinite equivalent codes, i.e. all those obtained by permutation or linear combination of the matrix G .




Fig. 14.2. Representation of the space generated by a code

Therefore, encoding a data message m consists in multiplying the message m by the code generator matrix G , according to Eq. (14.8).

$$c = m \cdot G \quad (14.8)$$

Definition 14.2 G is called in *standard form* or in *systematic form* if $G = (I_k, P)$, where I_k is the identity matrix $k \times k$ and P is a matrix $k \times (n - k)$. If G is in standard form then the first k symbols of a word are called information symbols.

Theorem 14.1 If a code $C[n,k]$ has a matrix $G = (I_k, P)$ in standard form, then there is a *parity matrix* $H = (-P^T, I_{n-k})$ where P^T is the transpose of P and it is a matrix $(n-k) \times k$ and I_{n-k} is the identity matrix $(n-k) \times (n-k)$.

Systematic codes have the advantage that the data message is visible in the codeword and can be read before decoding. For codes in non-systematic form the message is no more recognizable in the encoded sequence and it is necessary to have the inverse encoding function to recognize the data sequence.

Definition 14.3 The *code rate* is defined as the ratio between the number of information bits and the codeword length.

Definition 14.4 If C is a linear code with parity matrix H , then $x \cdot H^T$ is called *syndrome* of x .

All the codewords are characterized by a syndrome equal to 0.

The syndrome is the main actor of the decoding. Having received or read the message r , first of all it is necessary to understand if it is corrupted by calculating:

$$s = x \cdot H^T \quad (14.9)$$

There are two possibilities:

- $s = 0 \Rightarrow$ the message r is recognized as correct;
- $s \neq 0 \Rightarrow$ the received message contains some errors.

In this second case we use the maximum likelihood decoding procedure, that is we list all the codewords and we calculate the distance between r and the codewords. Therefore the vector c that has most likely been sent is the one which has the smallest distance from r .

Definition 14.5 It is called *minimum distance* or *Hamming distance* d of a code, the minimum number of different symbols between any two codewords.

We can see that for a linear code the minimum distance is equivalent to the minimum distance between all the codewords and the codeword 0.

Definition 14.6 A code has *detection capability* v if it is able to recognize all the messages, containing v errors at the most, as corrupted.

The detection capability is related to the minimum distance as described in Eq. (14.10).

$$v = d - 1 \quad (14.10)$$

Definition 14.7 A code has *correction capability* t if it is able to correct each combination of a number of errors equal to t at the most. The correction capability is calculated from the minimum distance d by the relation:

$$t = \left\lfloor \frac{d-1}{2} \right\rfloor \quad (14.11)$$

where the square brackets mean the floor function.

Between the length n of the codewords, the number k of information symbols and the number t of correctable errors, there is a relationship known as Hamming inequality:

$$\sum_{i=0}^t \binom{n}{i} \cdot (q-1)^i \leq q^{n-k} \quad (14.12)$$

Substantially, the number of parity symbols $(n - k)$ has to be sufficient to represent all the possible errors. A code is defined as *perfect* if the Eq. (14.12) is an equality. In this case all the vectors of the space are contained in circles of radius

$$t = \left\lceil \frac{d-1}{2} \right\rceil \quad (14.13)$$

around the codewords. Perfect codes have the highest information rate of all the codes with the same length and the same correction capability.

Codes are not fixed and they can be manipulated or combined depending on applications. The possible operation to increase the minimum distance of a code is the extension.

Definition 14.8 A code $C[n, k]$ is *extended* to a code $C'[n+1, k]$ by adding one more parity symbol.

Generally, for binary codes, the additional parity bit is the total parity of the message. This is calculated as *sum modulo 2* (XOR) of all the bits of the message. We can observe that, by itself, this type of code, called parity code, is enough to detect, but not to correct, an odd number of errors.

In a lot of applications there are external factors not subject to error check which determine the length permitted to an error correction code. Non volatile memories, for example, operate on codewords that have a length power of 2.

When the “natural” length of the code is not suitable it is possible to change it with the shortening operation.

Definition 14.9 A $C[n, k]$ is *shortened* into a code $C'[n-j, k-j]$ by erasing j columns of the parity matrix.

Observe that both the shortening operation and the extension are applicable only to linear codes.

Another operation that can be performed with codes is the concatenation (Fig. 14.3). This operation is based on the combination of two different codes called inner code (C_{IN}) and outer code (C_{OUT}). Concatenated codes were first developed in 1966 by Dave Forney.

Here is the concept. C_{IN} is decoded using a maximum-likelihood approach; in this way, the error probability decreases with increasing length. C_{OUT} is a block code of length N and, therefore, it can be decoded in polynomial time of N . The resulting concatenated code $C_{IN} \circ C_{OUT}$ combines together the error probability property of the inner code and the decoding time property of the outer code.

Typically, the inner code is a soft-decision convolutional code with a short length and the outer code is a longer hard-decision block code (e.g. BCH or Reed Solomon). As the erroneous output of the convolutional code is bursty, the outer code must be more robust to burst errors. Additionally, an interleaving layer may be used that spreads burst errors across a wider range.




Fig. 14.3. Schematic diagram of the codes concatenation

14.2.2 Convolutional codes

Convolutional codes have been widely used in applications like satellite communications. The greatest difference in comparison with block codes is that the encoder has a memory and this means that, at every clock pulse, the bits encoded depend not only on the k th information bit but also on the preceding m ones. Besides, a convolutional encoder converts the whole information string, without taking into account the length, into a codeword. In this way not all the codewords have the same length.

The construction of the convolutional codes is not based on the distance because it is not possible to define a distance for them. For this reason, there aren't closed formulas describing the error correcting capability of the code and it must be simulated. In Fig. 14.4 there is a comparison between the correction capabilities of convolutional and block codes.

In Fig. 14.5 an example of a convolutional code is represented. The squares are memory elements. The system is ruled by an external clock that produces a signal every t_0 seconds, the effect is that the signals move in the direction of the arrows toward the following element. The output therefore depends on the current input and on the two preceding ones.

Given an input $c = (\dots, c_{-1}, c_0, c_1, \dots, c_b, \dots)$ where l is a second of time, the outputs are two sequences v^1 and v^2 : $v^1 = (\dots, v_{-1}^1, v_0^1, v_1^1, \dots, v_l^1, \dots)$, $v^2 = (\dots, v_{-1}^2, v_0^2, v_1^2, \dots, v_l^2, \dots)$. At time l the input is c_l and the output is $v^l = (v_l^1, v_l^2)$ with:

$$v_l^1 = c_l + c_{l-2} \quad (14.14)$$

$$v_l^2 = c_l + c_{l-1} + c_{l-2} \quad (14.15)$$

where the sum is meant as binary addition.

Consequently the code is linear and with memory 2. The codeword is formed as:

$$v = (\dots, v_{-1}^1, v_{-1}^2, v_0^1, v_0^2, v_1^1, v_1^2, \dots, v_l^1, v_l^2, \dots) \quad (14.16)$$

As for the block codes, the encoding is calculated as $v=c \cdot G$. Called D the delay operator, v and c can be written as:

$$c(D) = \dots + c_{-1}D^{-1} + c_0 + c_1D + c_2D^2 + \dots \quad (14.17)$$

$$v^i(D) = \dots + v_{-1}^iD^{-1} + v_0^i + v_1^iD + v_2^iD^2 + \dots \quad (14.18)$$




Fig. 14.4. Probability of error occurrence and error correction

Sequence generators, also known as *generator polynomials*, can be written as polynomials using the same approach.

Now consider the code of Fig. 14.5. The generator polynomials are:

$$g_1 = (101) = 1 + D^2 \quad g_2 = (111) = 1 + D + D^2 \quad (14.19)$$

The decoding algorithm used in this kind of codes is the algorithm of maximum likelihood. The discoverer of this algorithm is Andrea Viterbi. The state of the encoder is defined as the content of memory cells. If K is the number of memory elements, the number of states is 2^K . The state diagram consists of knots, representing the states, and lines, representing the state transitions. Figure 14.6a shows the state diagram of the code introduced in Fig. 14.5.




Fig. 14.5. Example of convolutional code (2,1,2)

For example, if we are in the state S_2 and the input is a 0, we reach the state S_1 , meanwhile the output, following Eqs. (14.14) and (14.15), will be 01. On the other hand, if the input is a 1, we can reach the state S_3 and the output will be 10.

A Trellis diagram [2, 3] can be deduced from the state diagram by tracing all the possible input/output sequences and the state transitions (Fig. 14.6b).

The Viterbi decoding algorithm makes use of the Trellis diagram and describes a procedure of maximum likelihood decoding. The likelihood is described through a *metric* like the Hamming distance (Definition 14.5). Of course, different metrics are available and the designer has to choose the most suitable depending on the transmission channel.

Viterbi Algorithm is made up by five steps as describe below.

- Step 1: since it is not possible to wait for the whole word which could be endless, truncate the word at time b .
- Step 2: beginning at time $t = 0$, compute the partial metric for the single path entering each state. Store the survivor path and its metric for each state.
- Step 3: increase t . Compute the partial metric for all the 2^K paths entering a state by adding the branch metric entering that state to the metric of the connecting survivor at the previous unit.
- Step 4: for each state, compare the metrics of all the 2^K paths entering that state, select the path with the smallest metric (the survivor) and delete all other paths.
- Step 5: if $t < b$ repeat steps 3 and 4, otherwise stop.

The decoding complexity of this algorithm grows exponentially with the number of memory elements used in the convolutional encoder. The disadvantage in the use of convolutional codes is that a constructive method to generate them does not exist; for values K of practical interest the search for good codes must be carried out in an exhaustive way by trying out all the possible connections, with some rules to reject the bad solutions as fast as possible.

Besides, despite the sure advantage of a high hardware parallelism, the decoding requires a lot of memory to store the surviving paths and it is necessary to wait for the end of the transmission to decide on all the information bits. Their greatest use is in concatenated codes and in turbo codes, generally combined with a cyclic code or with a non-symmetric error channel.




Fig. 14.6. (a) State Diagram and (b) Trellis diagram of circuit in Fig. 14.5

14.3 BCH codes

BCH codes belong to the most important class of cyclic algebraic codes. They were found through independent researches by Hocquenghem in 1959 and by Bose and Ray-Chauduri in 1960 [5, 6].

For BCH codes the minimum distance can be ensured during construction. Generally speaking, in order to know the minimum distance for a linear code with generator polynomial $g(x)$, it is necessary to compute the distance between all the possible codewords. BCH codes, by imposing some constraints on the generator polynomial, are able to ensure a "designed distance".

Definition 14.10 Let β be an element of the Galois Field $GF(q^m)$. Let b be a non-negative integer. A BCH code with "designed" distance d is generated by the polynomial $g(x)$ of minimal degree that has $d - 1$ consecutive powers of β : $\beta^b, \beta^{b+1}, \dots, \beta^{b+d-2}$ as roots. Given ψ_i the minimal polynomial of β^{b+i} for $0 \leq i < d - 1$, $g(x)$ is computed as:

$$g(x) = LCM \{ \psi_0(x), \psi_1(x), \dots, \psi_{d-2}(x) \} \quad (14.20)$$

and the data carried by the code is $k = n - \deg(g(x))$.

It is possible to show that the designed d is at least $2t + 1$, hence the code is able to correct t errors.

If we assume $b=1$, and β a primitive element of $GF(q^m)$ the code becomes a *narrow-sense* and *primitive* BCH code of length $q^m - 1$ able to correct t errors. We shall now consider narrow-sense primitive BCH codes.

In general, the decoding of a BCH code is at least 10 times more complicated than the encoding. In the following we will deal only with binary BCH codes, whose structure is presented in Fig. 14.7.




Fig. 14.7. General structure of a binary BCH code

14.3.1 Encoding

Suppose we have a BCH code $[n, k]$ with generator polynomial $g(x)$ and a message $m(x)$ to encode, written as a polynomial of degree $k - 1$.

First of all, the message $m(x)$ is multiplied by x^{n-k} and subsequently divided by $g(x)$, obtaining a quotient $q(x)$ and a remainder $r(x)$ in accordance with Eqs. (14.21) and (14.22).

$$\frac{m(x) \cdot x^{n-k}}{g(x)} = q(x) + \frac{r(x)}{g(x)} \quad (14.21)$$

$$m(x) \cdot x^{n-k} + r(x) = q(x) \cdot g(x) \quad (14.22)$$

The multiplication of the message $m(x)$ by x^{n-k} produces, as a result, a polynomial of degree $n - 1$ where the first $n - k$ coefficients, now null, will then be occupied by parity bits.

Therefore the encoded word $c(x)$ is calculated as:

$$c(x) = m(x) \cdot x^{n-k} + r(x) \quad (14.23)$$

Practical implementation of Eq. (14.23) is depicted in Fig. 14.8. Note that, since we are considering binary BCH codes, the sum is actually a XOR, while the product is an AND.

The structure is formed by $n - k$ registers and there can be a maximum number of $n - k$ multiplications by the coefficients of the generator polynomial. The content of the registers has to be summed to the incoming bit if the corresponding

coefficient of the generator polynomial is 1, otherwise there will not be any wire indicating the sum.

For the first k steps, the message inputs the structure and contemporarily goes out to form the codeword. At the end, the $n - k$ registers contain the parity bits. The parity bits output in $n - k$ clock pulses. Note that the latency of this structure is n .




Fig. 14.8. Binary BCH encoder

14.3.2 Decoding

The decoding operation follows three fundamental steps, as shown in Fig. 14.7.

- Calculation of the syndromes
- Calculation of the coefficients of the error locator polynomial (usually done with Berlekamp–Massey algorithm [1, 6–8])
- Calculation of the roots of the error locator polynomial (usually done with Chien algorithm [1, 2, 4])

During the transmission or the reading of the encoded message some errors may occur. Also the errors can be represented by a polynomial that has coefficient 1 in correspondence with every error's position:

$$E(x) = E_0 + E_1x + \dots + E_{n-1}x^{n-1} \quad (14.24)$$

Observe that, in order for the code to be corrector of t errors, at most t non-null coefficients are allowed in Eq. (14.24).

The read vector $R(x)$ is therefore:

$$R(x) = c(x) + E(x) \quad (14.25)$$

The first decoding step consists in calculating the $2t$ syndromes for the read message:

$$\frac{R(x)}{\psi_i(x)} = Q_i(x) + \frac{S_i(x)}{\psi_i(x)} \quad \text{with } 1 \leq i \leq 2t \quad (14.26)$$

$$S_i(x) = Q_i(x) \cdot \psi_i(x) + R(x) \quad \text{with } 1 \leq i \leq 2t \quad (14.27)$$

In accordance with Eqs. (14.26) and (14.27), the received vector is divided by each minimal polynomial ψ_i forming the generator polynomial, thus getting a quotient $Q_i(x)$ and a remainder $S_i(x)$ called *syndrome*.

At this point the $2t$ syndromes must be evaluated into the elements $\beta, \beta^2, \beta^3, \dots, \beta^{2t}$ whose ψ_i are the minimal polynomials. With reference to Eq. (14.28) this evaluation is reduced to an evaluation of the message received in $\beta, \beta^2, \beta^3, \dots, \beta^{2t}$, since $\psi_i(\beta^i) = 0$ (for $1 \leq i \leq 2t$) by definition of minimal polynomial.

$$S_i(\beta^i) = S_i = Q_i(\beta^i) \cdot \psi_i(\beta^i) + R(\beta^i) = R(\beta^i) \quad (14.28)$$

Consequently, the i th syndrome can be calculated either as a remainder of the division between the received message and the minimal polynomial ψ_i , then evaluated in β^i , or as the evaluation in β^i of the received message.

Observe that, in case no errors occur, the polynomial received is a codeword: therefore the remainder of the division of Eq. (14.26) is null and all the syndromes are identically null. On the other hand, verifying if the syndromes are identically null is a necessary and sufficient condition to understand if the read message is a codeword or if some errors occurred.

For binary codes we use the property:

$$S_{2i} = S_i^2 \quad (14.29)$$

It is therefore necessary to calculate only t syndromes. The syndromes are linked to the error positions, in fact:

$$S_i = R(\alpha^i) = c(\alpha^i) + E(\alpha^i) = E(\alpha^i) \quad (14.30)$$

By indicating the error positions with X and the number of errors that have really occurred with v (for a code corrector of t errors it must be $v \leq t$), we get Eq. (14.31).

$$\begin{aligned} S_i &= E(\alpha^i) = x^{e_1} + x^{e_2} + \dots + x^{e_v} \\ &= X_1^i + X_2^i + \dots + X_v^i = \sum_{l=1}^v X_l^i \end{aligned} \quad (14.31)$$

From Eq. (14.31) we get Eqs. (14.32) called *power-sum symmetric functions*:

$$\begin{aligned} S_1 &= X_1 + X_2 + \dots + X_v \\ S_3 &= X_1^3 + X_2^3 + \dots + X_v^3 \\ &\vdots \\ S_{2t-1} &= X_1^{2t-1} + X_2^{2t-1} + \dots + X_v^{2t-1} \end{aligned} \quad (14.32)$$

Each method for solving these equations is a BCH decoding procedure.

Definition 14.11 It is defined as *error locator polynomial* $\Lambda(x)$ the polynomial whose roots are the inverse of the error positions.

From the definition we have:

$$\Lambda(x) = \prod_{i=1}^v (1 - xX_i) \quad (14.33)$$

Observe that the degree of the error locator polynomial gives the number of errors that occurred. The degree of $\Lambda(x)$ is t at most, so, in the case more than t errors occur, the polynomial $\Lambda(x)$ could erroneously indicate t or less errors.

For each error X_j it must be:

$$1 + \Lambda_1 X_j^{-1} + \dots + \Lambda_{v-1} X_j^{-v+1} + \Lambda_v X_j^{-v} = 0 \quad (14.34)$$

Multiplying by X_j^{i+v} we get

$$X_j^{i+v} + \Lambda_1 X_j^{i+v-1} + \dots + \Lambda_{v-1} X_j^{i+1} + \Lambda_v X_j^i = 0 \quad (14.35)$$

then, by summing over j ,

$$\sum_{j=1}^v X_j^{i+v} + \Lambda_1 \sum_{j=1}^v X_j^{i+v-1} + \dots + \Lambda_v \sum_{j=1}^v X_j^i = 0 \quad (14.36)$$

$$S_{i+v} + \Lambda_1 S_{i+v-1} + \dots + \Lambda_v S_i = 0 \quad (14.37)$$

which rewritten in matricial form becomes:

$$\begin{pmatrix} S_{v+1} \\ S_{v+2} \\ S_{v+3} \\ \vdots \\ S_{2v-1} \end{pmatrix} = \begin{pmatrix} S_1 & S_2 & S_3 & \cdots & S_v \\ S_2 & S_3 & S_4 & \cdots & S_{v+1} \\ S_3 & S_4 & \cdots & \cdots & S_{v+2} \\ \vdots & \vdots & \vdots & & \vdots \\ S_v & S_{v+1} & S_{v+2} & \cdots & S_{2v-1} \end{pmatrix} \cdot \begin{pmatrix} \Lambda_v \\ \Lambda_{v-1} \\ \Lambda_{v-2} \\ \vdots \\ \Lambda_1 \end{pmatrix} \quad (14.38)$$

Consequently, once the syndromes are known, it is possible to find the coefficients of the error locator polynomial.

The last step of the decoding process consists in searching for the roots of the error locator polynomial. If the roots are separate and they are in the field, then it is enough to calculate their inverse to have the error positions. If they are not separate or they are not in the correct field, it means that the word received has a distance from a codeword greater than t . In this case an incorrigible error pattern occurred and the decoding process fails.

14.4 Reed–Solomon codes

Reed–Solomon codes were described for the first time in 1960, in the article “Polynomial codes over certain finite fields” by Reed and Solomon [9]. It is only later that the close link with BCH codes was understood.

There are a lot of ways to present Reed–Solomon codes; in the following we will prefer the one that more closely links them to BCH codes. In this way, we will be able to underline that Reed–Solomon codes are a subset of BCH codes q -ari.

Looking back to the Definition 14.10 of BCH codes we can affirm:

Definition 14.12 Reed–Solomon codes are BCH codes q^m -ari with $n=q^m - 1$.

The construction of the generator polynomial $g(x)$ results to be very simple, as the minimal polynomial in $GF(q^m)$ of an element β in $GF(q^m)$ is simply $x - \beta$. Generally Reed–Solomon codes correct symbols instead of bits.

Therefore, a generator polynomial of a Reed–Solomon code has the form

$$g(x) = (x - \beta^b) \cdot (x - \beta^{b+1}) \cdots (x - \beta^{b+2t-1}) \quad (14.39)$$

and the code has $k=q^m - 1 - 2t$ information symbols.

Similarly to BCH codes, assuming $b = 1$ we get narrow-sense codes, even though, in applications, the choice $b = 0$ often simplifies the computational complexity.

A particular case, and surely the most used in applications, concerns Reed–Solomon codes in $GF(2^m)$, where each symbol is composed by m bits.

Suppose, for example, we search for a code whose symbol is a byte, namely 8 bits. By definition of Reed–Solomon codes, the codeword length is fixed at $2^8 - 1 = 255$ bytes. In the particular case of a code able to correct two symbols, we need four parity bytes.

Now, suppose we wish to construct a Reed–Solomon code in $GF(8)$ corrector of two error symbols. The primitive polynomial of the field is $x^3 + x + 1$. The code is seven symbols long, three of them are information symbols. The generator polynomial must have four roots in $GF(8)$:

$$\begin{aligned} g(x) &= (x - \alpha) \cdot (x - \alpha^2) \cdot (x - \alpha^3) \cdot (x - \alpha^4) \\ &= x^4 + \alpha^3 x^3 + x^2 + \alpha x + \alpha^3 \end{aligned} \quad (14.40)$$

The code just described corrects two symbols, each composed by 3 bits. It is however wrong to deduce that such code corrects 6 bits, since, if errors occur in

different symbols, the code is unable to correct them. Thus, the code is able to correct a number of bits varying from 2 to 6 depending on the position of the errors inside the symbols. Due to the fact that they were born to correct symbols, Reed–Solomon codes are particularly suitable when burst type errors occur.

The general structure of a Reed–Solomon code is that of Fig. 14.9. Reed–Solomon codes require a further step (Forney algorithm [1–4]) during the decoding process, compared to binary BCH codes, since it is not enough to know the error positions but it is also necessary to know their magnitude.




Fig. 14.9. Structure of a Reed–Solomon code

14.5 BCH versus Reed–Solomon

The choice of the more suitable cyclic code for a particular application is not simple and it would be useful to know the distribution and the generation of the errors in details.

For this reason, it is not possible to say whether a Reed–Solomon code is superior or inferior to a binary BCH code. The greatest difference between the two codes is the correction basic unit, as BCH code corrects bits, while the Reed–Solomon code corrects symbols. Therefore, a BCH code corrects t bit errors, while the Reed–Solomon code corrects t symbols of s bit errors, which can be a variable number of bits: from t bits if errors occur in different symbols up to st bits in case errors occur consecutively. At a first glance this may seem an advantage, which however is fully exploited only if the error source is of burst type which is not the case for NAND memories (Chaps. 13 and 4).

The choice of a code can also be based on the occupied area. In this case we have to note that both hardware implementation and the number of parity bits contribute to increase area. Regarding this last aspect, fixing the same correction capability, Reed–Solomon codes require a greater number of parity bits, as shown in Table 14.1.

For what concerns the encoding implementation, the difference is surely that binary codes require binary operators (AND and XOR), while Reed–Solomon codes require operators in $GF(2^m)$ which have a greater computational complexity.

Besides, since Reed–Solomon codes also require a greater number of parity bits, also the generator polynomial has a higher degree, and consequently the number of operators results to be higher. The remainders of the divisions must be stored into registers. Given the structure of the two codes, while BCH codes require the storage of bits, for Reed–Solomon codes it is necessary to store symbols.

Table 14.1. Number of parity bits required by Hamming inequality (minimum), BCH and RS codes

Data bits	<i>Min</i> $T = 1$	<i>BCH</i> $t = 1$	<i>RS</i> $t = 1$	<i>Min</i> $t = 2$	<i>BCH</i> $t = 2$	<i>RS</i> $t = 2$	<i>Min</i> $t = 3$	<i>BCH</i> $t = 3$	<i>RS</i> $t = 3$	<i>Min</i> $t = 4$	<i>BCH</i> $t = 4$	<i>RS</i> $t = 4$
2,048	12	12	18	22	24	36	31	36	54	40	48	72
4,096	13	13	18	24	26	36	34	39	54	44	52	72
8,192	14	14	20	26	28	40	37	42	60	48	56	80
16,384	15	15	22	28	30	44	40	45	66	52	60	88
32,768	16	16	24	30	32	48	43	48	72	56	64	96

The hardware structure of the syndromes calculation involves the same considerations made for the encoding: the Reed–Solomon codes require a greater number of operators with a greater complexity. Besides, Reed–Solomon codes require the calculation of $2t$ syndromes, a number doubled if compared to binary BCH codes which exploit the binary property:

$$S_{2t} = S_t^2 \quad (14.41)$$

Consequently also the inputs to the Berlekamp algorithm are doubled in comparison to a BCH code. If the iterative Berlekamp algorithm [1–4] is chosen, BCH codes require t iterations instead of the $2t$ required by Reed–Solomon codes.

The same considerations made for the encoding implementation apply to the Chien algorithm. Despite the fact that the number of operators is not actually greater, their complexity, as well as the registers for the storage, are more onerous in the case of Reed–Solomon codes compared to BCH ones.

Finally, note that Reed–Solomon codes require a further step in the decoding, represented by the Forney algorithm. The most onerous operation in this step is the inversion, which is not simple and it requires an additional ROM in order to be efficient.

For what said so far, it is evident that if the choice criterion is the occupied area, BCH codes are by far superior to Reed–Solomon codes.

When the choice criterion becomes the latency then the situation changes and the BCH codes turn out to have inferior performances compared to Reed–Solomon codes. The difference lies in the fact that BCH codes are serial and process 1 bit at a time, while Reed–Solomon codes, operating on symbols, are defined with a parallel structure.

In conclusion, Reed–Solomon codes generally have a greater correction capability for burst type errors, but they require a bigger area. On the contrary, BCH codes

have a smaller area but a higher latency time. It is also possible to upgrade the BCH performances using a parallel architecture as shown in the next section.

Usually, NAND memories don't show a burst behavior of the failed cells; therefore, BCH is the preferred choice by the NAND users.

14.6 Parallel BCH

In this section we assume to have a system formed by a control unit which provides and receives data from a NAND memory on an 8-bit bus.

The BCH coding and decoding operations are described in the vast majority of the cases as a process which operates in a serial way on a sequence of bits. Hence the difficulty to adopt it in a system that operates on words. The solution consists in defining machines able to elaborate the received data in a parallel way [10–13].

The serial structure of the parity computation machine can be turned into a parallel machine following some simple steps. The iterative operation carried out by the serial divisor to elaborate a word, can be unrolled as shown in Fig. 14.10. In this way, the following state of the registers is calculated on the basis of all the operations needed to elaborate a whole message word.




Fig. 14.10. Description of the sequential implementation of the divider and of how it could be unrolled to get a parallel implementation

The high voltage section of the device (comprising charge pumps and regulators) has to provide precise voltages necessary for reliable array operations during the whole device's life.




Fig. 15.1. NAND Flash memory block diagram

If we consider a write/erase cycling specification for one block of the array equal to 10 k cycles, we must be aware that the charge pumps, providing the voltages necessary for the cycling, operate in the worst (unlikely) case for 10 k cycles multiplied by the number of blocks in the array. Even if those reliability aspects are taken into account during the design phase, the manufacturer must guarantee that the production is able to maintain the reliability targets from lot to lot.

Moreover, a control logic section is used to manage all the circuits involved in memory operations. This section is frequently made up by a microcontroller plus a mask ROM and a RAM for executing code and storing data. Other implementations can be based on programmable logic arrays [3]. The digital circuitry in the chip has the same testability issues as in digital ASICs; testability can be secured by proper testing tools and accessibility such as digital scan chains in order to enable the highest quality levels and minimum testing time.

The aim of the following sections is to provide an introductory view of the test issues related to array architecture and peripheral circuits. For detailed descriptions of the circuits, it is recommended to refer to specific chapters in this book and bibliography. Specific test mode descriptions will be given in Sect. 15.4.

In fact, as is the case for a 4LC device, the highest verification level must be low enough to prevent bit failures such as program disturb and read disturb. The more states a memory cell must be able to store, the more finely divided its program threshold window is. Therefore, it is mandatory to reduce further the width of programmed distributions (Sect. 16.2).

A reduction in area occupation of the matrix has a drawback in an increment of the area occupied by the peripheral circuits: in particular, the area of the sensing circuits. In fact, the number of latches inside a sensing circuit is (as a minimum) equal to the number n of bits stored inside the cell (Sects. 16.3 and 16.4).

A number n of bits corresponds to a number $(2^n - 1)$ of levels of read and verify (see Fig. 16.2). Therefore, a device with n bits per cell needs $(2^n - 1)$ SROs to read the content of the cell. Moreover, the number of verify operations after a single program pulse is equal to $(2^n - 1)$.

Taking into account the last item, we could come to the conclusion that a read operation of a 16LC device is 15 times slower than a SLC device. Indeed, 15 SROs can read four pages, and therefore the *average page access time* T_R

$$T_R = \frac{(2^n - 1)}{n} \cdot T_{SRO} \quad (16.1)$$

is equal to $3.75 T_{SRO}$, where T_{SRO} is the time spent by one SRO.

Therefore, T_R ratio between a 16LC and a SLC device is 3.75 and not 15. This ratio is as low as 2.5 if a 4LC device is considered ($T_R = 1.5 \cdot T_{SRO}$).




Fig. 16.2. 16LC ISPP and its 15 verify levels

The *CRC (Cyclic Redundancy Check)* is an error detection code. It checks the integrity of the command, which could be received incorrectly because of transmission errors. The CRC is computed over the command token. Upon reception of the command it is computed independently by the card and compared with the one received in the command. If they match, the command is executed otherwise an error is signaled.

Typical examples of commands are the Read and Write Single Block commands. Through these commands, a Block (usually 512 byte long) is read from or written to the card. Please note that the card write command can write both 0 and 1 and not just 0 as in Flash memories. The Command Index for the Read Single Block command is 17, while the one for the Write Single Block is 24. In both cases, the Argument is the data address to be read or written. The Argument field of just 32 bits should not mislead into believing that only $2^{32} = 4$ Gbyte could be addressed, because for card with a capacity greater than 2 Gbytes the addressing is on a sector basis. Consequently, any card density can be accessed with a 32-bit field argument.

The Read and Write Single Block commands are obviously followed by the data transfer, from the host to the card for the Write Block and from the card to the host for the Read Block.

Each command is responded to by the card. The response serves the purpose for the host to have feedback information about the command just issued, like whether the card has correctly received the command, has executed or aborted it etc. Sometimes the response itself contains the information requested by the host through the command. There are a few type of responses, the most common having a code length of 48 bits and the following structure:

	Start Bit	Transmission Bit	Command Index	Card Status	CRC7	End Bit
Width (bits)	1	1	6	32	7	1

The fields are clearly similar to the command token. Here the Command Index is the replica of the command being responded to. The Card Status is a card register stashing a lot of useful information about the card status and the outcome of the command being executed.

We can now delineate the whole behavior of the host-card communication during a Read and Write operation.

In a Read Single Block, the host issues the command on the CMD line (Fig. 17.3). The card sends its response indicating in the Status Register whether the command has been accepted or declined for any reason (address out of range, CRC failed...).

Upon reception of the command, the card starts retrieving the requested data. As soon as the data are fetched and ready to be transmitted, the card sends them via the DAT line.

The *Driver Stage Register* selects the power strength of the output buffer. This permits to adapt the card’s output stage to the transfer rate and to the bus capacitance (which varies for instance according to the bus length and the number of connected cards).

17.4 Flash translation layer

Operating systems usually write data in sectors (e.g. 512 bytes in size). Unfortunately, Flash memories lack the so-called *update-in-place* capability. This means they cannot be altered on a sector (page in NAND nomenclature) basis, because the erase operation involves a whole block (made up of several sectors/pages). Therefore, file systems cannot write ones back to zeros in a single addressed sector. This is awkward for a high-level file system to manage.

In order to alleviate the workload of the operating system, software routines can be implemented to emulate a NAND Flash behavior where single sectors (or NAND pages) can be modified independently. Such software is called *Flash Translation Layer (FTL)*, which is an intermediate layer between the file system and the storage media (the NAND Flash). The operating system can then write the NAND memory on a page basis without worrying about the details of its physical implementation. Usually it is the Controller in Fig. 17.1 taking care of the FTL.

The FTL achieves its task through a *logical to physical re-mapping*. Hence, when the file system sends the command to write or update a certain logical page, the FTL actually writes it in a different available physical page and updates a table storing the logical to physical mapping (*mapping table*). The page containing the older data is marked as invalid; Fig. 17.6 shows an example.




Fig. 17.6. Logical to physical re-mapping

Here page 0 has been written first at address n , followed by page 1. Then page 0 has been updated, but instead of overwriting it at address n , the FTL writes it in another location ($n + 2$) and marks address n as invalid.

As it is apparent from Fig. 17.6, after a while a great portion of the memory will be filled with invalid data taking precious unusable space. Eventually, all these locations will have to be erased to make this memory space available to new data. This procedure is not trivial, because the erase operation involves a whole block, comprising several pages, some with valid data, while others with invalid data. Therefore, valid data must be copied first to a new location, before the erase operation takes place. All this procedure is called *garbage collection*. An example is shown in Fig. 17.7.

All valid pages are copied into new locations (the order does not matter because the mapping table takes care of the correct correspondence between the logical and physical address). Then all the memory locations containing the old data (valid and invalid) are erased, thus creating space accessible to new data.

Garbage collection is very critical for many reasons. Let us start to better illustrate this point by introducing the notion of garbage collection efficiency.

Garbage collection efficiency is defined as the ratio of the number of invalid pages in a block to the total number of pages in that block. The higher the invalid pages within a block, the higher the efficiency. Indeed, when the number of invalid pages within a block is high we have a twofold benefit. First, only a few pages must be copied into new locations, thus limiting the program time and reducing the cycling efforts (greater lifetime) of the memory cells. Moreover, when erasing such block, a lot of space will be available for new data, because only few data were valid and still taking memory space.




Fig. 17.7. Garbage collection

When should the garbage collection be run? A high program throughput calls for a lot of free space. If the host requires more space than readily available, it will have to halt its operation and wait for the garbage collection to kick in and erase enough invalid pages. This will limit the sustainable program throughput. So apparently, the garbage collection should be run as often as possible. However, this in turn will limit the garbage collection efficiency. Therefore, an accurate trade-off will have to be determined to establish the best performance of garbage collection.

Flash memories' blocks can only be programmed and erased a limited number of times. This is the reason why a good FTL should also implement a *wear leveling* algorithm. A wear leveling software module makes sure all NAND blocks are cycled an identical number of times, avoiding some blocks ending their cycle life while others being cycled only a few times. Were not a wear leveling algorithm present, blocks with frequently changed data would endure much more cycles than blocks containing long-lived data.

There are two levels of wear leveling. In the *first level wear leveling* new data is programmed into a free block with the fewest write/erase cycles. Two techniques to achieve this are the chain method and the array method.

In the *chain method*, the pool of available blocks is organized in a chain (Fig. 17.8).

A simple round robin algorithm selects one block after another. When a block is erased, it is added to the chain.

In the *array method* an age vector stores the number of each block's program/erase cycles. This information can be used by the wear leveling algorithm to ensure a uniform exploitation of all the blocks.

In the *second level wear leveling* from time to time blocks containing long-lived data are moved (copied) into other blocks even though they never received an update command. The original block can then be used to store new incoming data. The second level wear leveling is triggered when the difference between the maximum and the minimum number of cycles per block reaches a predetermined threshold.

Devising an efficient FTL is a taxing job. Two main concerns are the clustering and the cleaning policy. A few examples of both will be reported, but let start off by introducing the concepts of hot and cold data.

Hot data are the most frequently updated data. *Cold data* (or non-hot data) are seldom updated data. In a given block, it is advantageous to gather (cluster) all hot data or all cold data, because blocks with hot data will soon become garbage. Cleaning such blocks is beneficial because they have high garbage collection efficiency.

How to effectively cluster hot data then? One method is the *Dynamic dAta Clustering (DAC)* depicted in Fig. 17.9.



Fig. 17.8. Chain method




Fig. 17.9. Dynamic dAta Clustering

In the DAC method, the Flash memory is divided into n regions, whose size is not necessarily identical. Each region is not even necessarily made up of contiguous memory locations. Region 1 contains the coldest data, Region n the hottest. The hotter a block becomes, the more it shifts towards Region n . Conversely, the colder a block becomes, the more it shifts towards Region 1. When a block is written for the first time, it is put in Region 1. When a block is updated, if it has been into its Region k for a time shorter than a $t_{\text{threshold}}$ it is written in Region $k + 1$, otherwise it is written again in Region k . The $t_{\text{threshold}}$ is introduced because a block's hotness degrades as it ages. Upon the cleaning request of a block, if its valid data have been in their Region k longer than $t_{\text{threshold}-2}$ they are copied in Region $k - 1$, otherwise they are still copied in Region k . Each block has associated with it the number of the Region it belongs to.

The real advantage of this technique is twofold. Firstly, the algorithm does not require complex calculations and secondly the classification is fine grained (which in turn improves the performance).

Another clustering method is the *Dynamic Striping*, illustrated in Fig. 17.10. In the Dynamic Striping methodology, two registers called LRU (least recently used) are utilized. Only logical block addresses belonging to the hot list are hot, all the others are cold. The candidate list contains the logical block addresses recently written. If a logical block address belongs to the candidate list and is updated after a short while, it is promoted to the hot list. The last element of the hot list is demoted to the candidate list.

As a further improvement of the Dynamic Striping, hot data are written in the bank with the lowest number of erase cycles to enhance the wear leveling.




Fig. 17.10. Dynamic striping

Cold data are written in the bank with the lowest capacity utilization (i.e. the ratio of the number of live pages to the number of total pages) in order to make the capacity utilization more uniform (cold data will remain longer in their locations).

Cleaning policies relate to the criterion used to pick up the block to be erased for garbage collection purposes.

The *greedy policy* is one of the simplest cleaning policies, in that it selects for erasing the block with the greatest number of invalid data. Though simple, this policy is not very efficient because it does not take into account the cycling the blocks have been subjected to. It has also been proved ineffective in case of localities of reference (i.e. the fact that if a certain memory location is referenced at a given time, it is probable that the neighboring memory locations will be referenced in the near future).

A better cleaning policy is the *cost-benefit policy*. The block chosen for erase is the one with the greatest value of the following expression:

$$\frac{age \cdot (1 - valid)}{2 \cdot valid} \quad (17.1)$$

where *valid* is the percentage of valid data in the block, and *age* is the time elapsed since the last modification. Note that $(1 - valid)$ is the percentage of invalid data, i.e. the space that ultimately will be freed. The term $2 \cdot valid$ can be conceptually regarded as the cost of reading the valid data and to copy them in another location ($2 \cdot valid = valid + valid$ and the first term is proportional to the time required to read the valid data while the second term is proportional to the time required to copy them in a new location). Therefore, the ratio $(1 - valid)/2 \cdot valid$ in the above expression can be considered as a benefit over cost ratio. This policy is better than greedy's because it takes into account the age of the block.

Yet another cleaning policy is the *Cost Age Time (CAT) policy*. The block chosen for erase is the one with the minimum value of the following expression:

$$\frac{valid}{1 - valid} \cdot \frac{1}{age} \cdot n_C \quad (17.2)$$

where *valid* and *age* are defined as in the previous case and n_C is the number of cleaning, i.e. the number of times a block has been erased. This cleaning policy is endowed even with a sort of wear leveling thanks to the term n_C .

A good FTL must also take care of *bad blocks*. Bad blocks contain not-working or unreliable bits and therefore cannot be used. They may be present since the manufacturing of the device or may develop during its lifetime. The FTL must store all the necessary information about which blocks can be used and which cannot. When a program operation in a certain block fails, the FTL acknowledges that through the NAND Status Register and acts copying the valid data of the block into a new good block; the block where the fail occurred is then be marked as bad.

Finally, a FTL must ensure the reliability of data in the event of sudden power loss. In order to guarantee this, the FTL must wait for the program operation to

complete successfully before marking the old page as invalid or before erasing a block. In such a way, if the new data is lost due to a power failure, the old data can still be recovered.

17.5 Cryptography

Security has become an important part of cards' characteristics. New protocols implement commands handling confidential data with the utmost caution. Let us mention the standard JESD84-A44 (eMMC), which states that every card complying with it must implement a special memory area called *Replay Protected Memory Block (RPMB)*, which must be protected against replay attacks. In a *replay attack* a valid data transmission is intercepted first, recorded and then played back again later on by an attacker. Hence, cryptography has become fundamental even in card design and usage. To protect a portion of the memory against replay attacks the eMMC protocol employs a key (MAC) and a nonce. To better understand these concepts, let us briefly review a few pivotal definitions and notions about cryptography.

A *cipher* is an algorithm to *encrypt* and *decrypt* a *plaintext* into and from a *ciphertext* (Fig. 17.11).




Fig. 17.11. Cipher

Note that in the everyday vernacular language sometimes the terms cipher and code are used interchangeably. However, in cryptography a *code* has a *codebook* associating each word or phrase of the plaintext with one or more codewords. Therefore, the plaintext set is somewhat limited and must be anticipated a priori. Codes operate according to the meaning while ciphers operate according to a single letter or bit.

Block ciphers work on blocks whereas *stream ciphers* work on continuous streams of data. Each cipher needs a *key* or *cryptovalue*. *Symmetric key algorithms* have the same key for encryption and decryption. *Asymmetric key algorithms* have different keys for encryption and decryption. In symmetric key algorithms, the key must be known only by the sender and by the recipient.

A *cryptographic hash function* is a function mapping a *message* into a message *digest* (a.k.a. *hash value*) as depicted in Fig. 17.12.

The memory, in which user data can be stored and the added memory necessary to implement the code, can be managed through the microcontroller, as shown in Fig. 17.21.

The microcontroller, apart from managing the two types of memories (they can be of different types and therefore having dedicated interfaces), is also requested to manage coding and decoding of the selected code.

The memory which stores the added information of the code, “Auxiliary Memory” in Fig. 17.21, can be used with different methods. A first approach is to evaluate the error correction capability needed for the total size of the memory we need to protect, the minimum codeword length (chunk), type and capability of the error correction code. Once the choices are made, it is possible to find the minimal additional memory compatible with the requests.

Using this method it is possible to link a chunk of the memory to be protected to a portion of the additional memory; this solution allows a fixed error correction capability for every chunk.

A second approach is based on a “dynamic” error correction capability (i.e. the error correction capability for a single chunk is increased when an error occurs); in this way it is possible to read correctly even in the case of further errors.

Specification to be fixed for implementation are, beyond of the size of the additional memory, the minimum and the maximum error correction capabilities (N_{min} and N_{max}).

At the beginning, every chunk of the primary memory is linked to a portion of the “Auxiliary Memory” needed to correct a minimum number of errors. Multiplying the total number of chunks by the parity bits needed to correct N_{min} errors, we can calculate the required portion of the memory.




Fig. 17.21. Auxiliary memory scheme used to store code bits

Table 17.2. NAND Flash: SLC versus MLC


	SLC	MLC
Sequential read access time	25 ns	25 ns
Random read access time	25 μ s	60 μ s
Program time	200 μ s	800 μ s
Erase time	2 ms	2.5 ms
Power supply – Vdd	1.8/3.0 V	3.0 V
Program/Erase cycles	100 k	10 k
Data retention	10 years	10 years

Many market applications require both SLC features to reliably store system code and data and MLC features, to save high density data while requiring lower reliability

A new solution, we'll call it flexible, is realized thanks to the microcontroller and consists on the partitioning of the single NAND cell in two sub sections, the first of them to be used as SLC, the other one as MLC (see Fig. 17.23).

The partitioning allows the microcontroller to use different algorithms and settings in memory usage depending on the features of the selected portion. The partition is dynamic; the space assigned to the code can be modified by the microcontroller. In the SLC partition program (PV) and read (RV) references are more relaxed and the pulses can have a wider size due to a single distribution in the positive side of the threshold voltage values. This minor precision in Vth values allows speeding up program and read operations, getting closer to native SLC memories. Additionally, in SLC partition, there is not the need of a high ECC capability, affecting as well performances and data area to store parity bits.

The same approach can be applied to 3 or more bits per cell devices, where a portion is reserved for a SLC-mode usage to get better performances in terms of throughput and reliability.

**Fig. 17.23.** Flexible solution scheme

18.2 Analysis of SSD performance

As shown in Fig. 18.2, SSD contains more than 8 NAND chips. For example, in Fig. 18.2, each NAND Flash memory is assumed to have 16 Gbit or 2 GByte capacity. One SSD is composed of 64 NAND chips and the total capacity is 128 GBytes, which is most suitable for the lap top computers. In Fig. 18.2, 4-channel configuration is adopted. The NAND controller can operate four channels independently.




Fig. 18.2. Hardware architecture of SSD

The NAND controller can issue a write command to four channels, writing four NAND chips at the same time and improving the system-level write performance. In a channel, 16 NAND Flash memory chips are connected where I/O signals and control signals such as the WE (Write Enable)-signal, the RE (Read Enable)-signal, the CLE (Command Latch Enable)-signal, and the ALE (Address Latch Enable)-signal are shared to save area for the bonding and the packaging. To select one NAND chip in a channel, different CE (Chip Enable)-signals and R/B (Ready/Busy)-signals are assigned to the NAND chips in a channel.

Figure 18.3 shows the software architecture of SSD. The NAND controller is composed of the host interface, the Flash translation layer (FTL) and the NAND interface. FTL is the central part of the NAND controller and perform a bad block management, a logical-physical address translation, a wear-leveling, ECC and an interleaving. The intelligent write algorithm described in Sects. 18.5 and 18.6 is also controlled with FTL.




Fig. 18.3. Software architecture of SSD

In a NAND Flash memory, the programming is slower than the read by one order of magnitude [1]. The typical random read time of a 2 bit/cell NAND Flash memory is 50 μ s while the typical random program time is 800 μ s. Considering the read/write access time of HDD is a few milliseconds, the key design challenge of SSD is to improve the write performance. To accelerate the write speed of SSD, an interleaving is proposed where multiple NAND chips in SSD are programmed at the same time [5]. The write speed of SSD with an interleaving is expressed as follows:

$$Performance_SSD = N \cdot Performance_NAND \quad (18.1)$$

N is the number of NAND Flash memories operated simultaneously. $Performance_NAND$ is the program speed of one NAND flash memory. As a memory cell is scaled down or more bits are stored in a memory cell such as 3 or 4 bit/cell, a more precise control of the memory cell V_{TH} is required and accelerating the NAND speed, $Performance_NAND$, becomes difficult. Actually, the program time of the 56 nm 3 bits/cell is 1.2 ms [6] and that of the 70 nm 4 bit/cell is 6 ms [7] which are much longer than that of the 2 bits/cell. Thus, the best strategy to improve the SSD speed is to increase the number of NAND chips in parallel, N . However, if N is maximized and all NAND chips in SSD, for example 64 NAND chips, are programmed simultaneously, 64 times as much current flows and the current consumption increases to an unacceptable high value, two Amperes. Therefore, in an actual SSD operation, the maximum number of N is restricted by the current consumption constraint.