

The Application of ECC/DSP to Flash Memory

*Idan Alrod, Dr. Eran Sharon, Dr. Idan Goldenberg, Ran Zamir, Alex Bazarsky,
Omer Fainzilber, Dr. Ishai Ilani, Ariel Navon, Stella Achtenberg & Dudy Avraham*

Western Digital Corporation

Contents

Introduction.....	3
System Design Challenges.....	3
ECC-Centered Solutions.....	5
RTC – Read Threshold Calibration	7
Memory Error Model (MEM) Estimator	8
Joint LDPC and RAID Decoding for Enterprise SSD	11
Additional ECC/DSP features.....	11
Summary.....	11
Appendix – Case Study	12

Table of Figures

Figure 1: Product requirements, NAND technology: Opposing trajectories.....	3
Figure 2: Variability of Cell Voltage Distributions: An Illustration	4
Figure 3: The need for handling high BER events.....	4
Figure 4: Generic structure for a flash-based system	5
Figure 5: Example of various decoding gears correction capability and overall performance	6
Figure 6: Cell voltage distribution estimate based on performing many reads.....	7
Figure 7: Example of RTC reads: TLC	8
Figure 8: Comparison of RTC vs. Valley Search, optimal results shown as reference... 8	8
Figure 9: Channel error models.....	9
Figure 10: Iterations between decoder and MEM estimator	9
Figure 11: Performance under a random channel with fixed mutual information	10
Figure 12: Read Thresholds.....	10
Figure 13: Tanner graph of a combined LDPC-RAID scheme.....	12
Figure 14: Combined LDPC-RAID failure reconstruction scheme.....	13
Figure 15: UBER curve for LDPC, optimal LDPC-XOR and simplified LDPC-XOR.....	13

Glossary

ASIC	Application Specific Integrated Circuit
BER	Bit Error Rate
BF	Bit-Flipping
CVD	Cell Voltage Distribution
CRC	Cyclical Redundancy Check
DSP	Digital Signal Processing
ECC	Error Correction Coding
LDPC	Low Density Parity Check
LLR	Log-Likelihood Ratio
MEM	Memory Error Model
MLC	Multi-level cell
MMD	Mismatch Memory Model
NAND	Not AND (represents a very dense architecture of Flash)
PRNG	Pseudo-Random Number Generation
QLC	Quad-level cell
RAID	Redundant Array of Independent Disks
RAM	Random Access Memory
RBER	Raw Bit Error Rate
RTC	Read Threshold Calibration
SSD	Solid State Drive
SW	Syndrome Weight
TLC	Triple-level cell
UBER	Unrecoverable Bit Error Rate
WDC	Western Digital Corporation
XOR	Exclusive OR

Introduction

Looking at NAND flash-based product requirements, we see they are derived from the market demand for higher memory capacities at lower costs which in turn results in a continuous increase in densities through process scaling, 3D stacking, and storage of more bits per memory cell (MLC → TLC → QLC) a.k.a. logical scaling. This trend causes fundamentally "noisier" media. Ensuring data integrity and reliable storage over such media requires protection against both random errors and memory defects.

In parallel to this density increase trend, new applications are driving requirements for higher performance, lower power consumption, and wider span of operational conditions (endurance, retention, cross temperature, etc.) in conjunction with higher reliability level and quality of service (QoS) requirements. Figure 1 below shows this disparity.

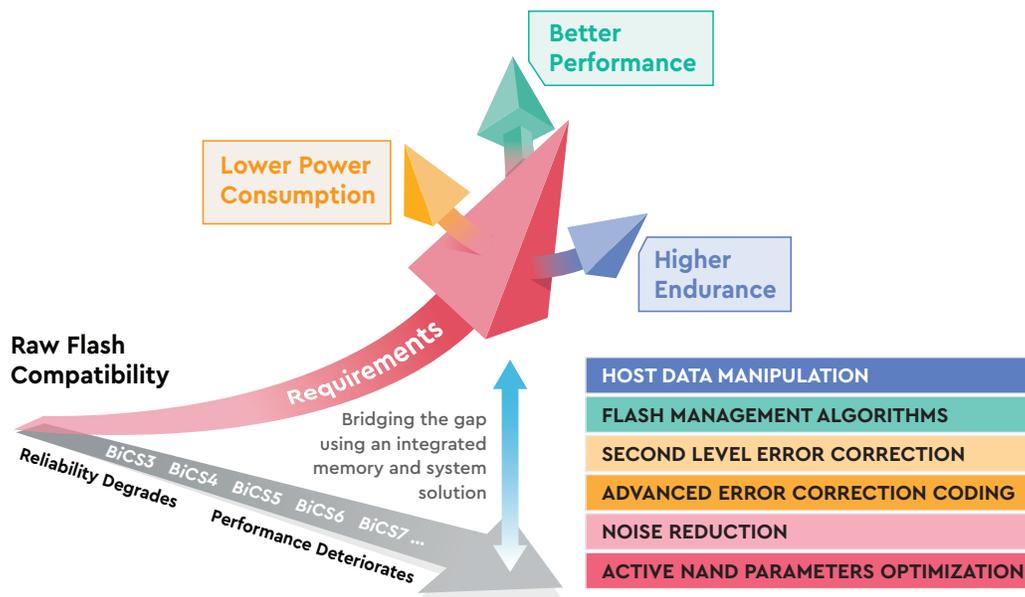


Figure 1: Product requirements, NAND technology: Opposing trajectories

In this white paper, we will explain why a vertically integrated memory and system solution is required in order to bridge the gap between the stringent requirements and the degraded media capabilities.

System Design Challenges

One of the main challenges due to the technology node shrinking and 3D stacking during the NAND fabrication process is maintaining process uniformity. This leads to increased variability between dies blocks and word-lines behavior at any point in time. The variability is exacerbated by the requirement to support a wide range of operational conditions, such as amount of P/E cycles, data retention times, cross temperature, and the ability to recover from a sudden voltage drop.

As an example of the effect of the variability, consider the cell voltage distributions of two different word-lines within the same memory, which were programmed at the same time, but are read at different temperature conditions, as shown in Figure 2 below. It clearly shows that the underlying distributions are significantly different between these two conditions.

The high variability leads to a wide BER distribution, which may be modeled as a log-normal distribution (See Figure 3 below). The distribution has low median BER (observed under most conditions) but a long upper tail that may reach high BER values (in outlier pages and under certain operational conditions).

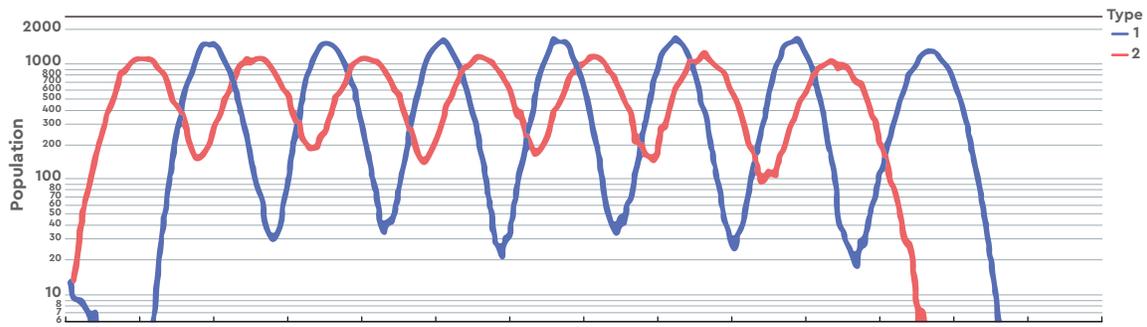


Figure 2: Variability of Cell Voltage Distributions: An Illustration

As memory systems have very stringent UBER specs (e.g., allowed UBER is $1e-16$), even if most pages under most conditions are exhibiting low BER, the system still needs to ensure the data integrity for the rare outlier pages and under extreme operational conditions as illustrated in Figure 3 below.

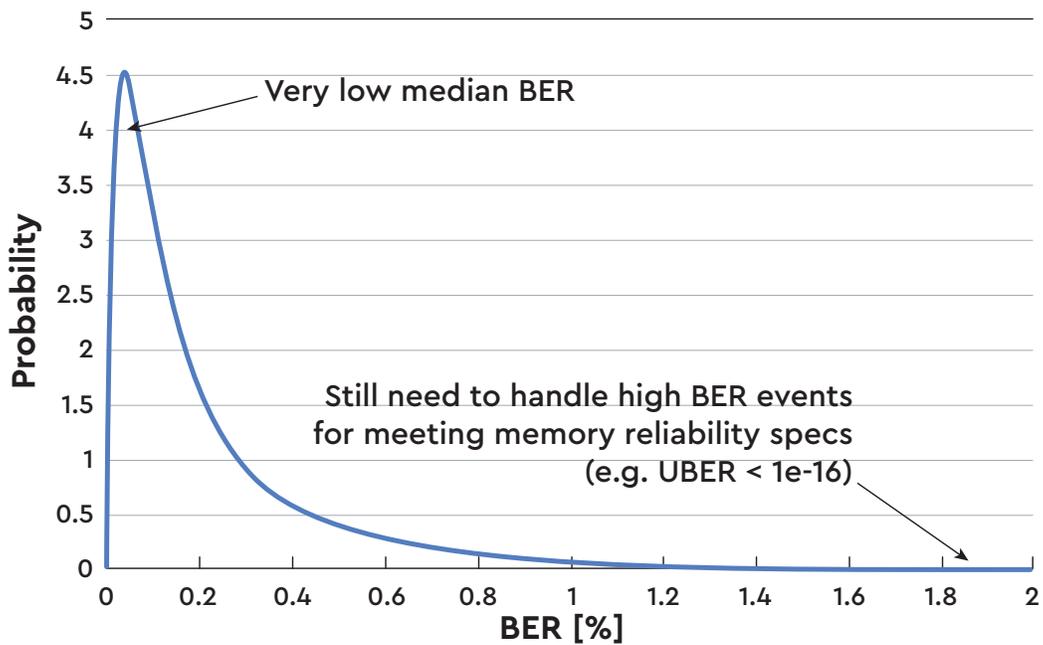


Figure 3: The need for handling high BER events

Following is a list of some of the tasks any memory management system needs to handle:

1. Read level tracking. The system must be able to actively monitor the cell voltage distribution and to track the optimal read threshold voltage levels. Scenarios such as the one depicted in Figure 2 demonstrate that this is an absolute necessity.
2. Wear leveling and health monitoring. Flash technology tends to degrade over time when the same physical page is repeatedly written and erased. To maximize the amount of possible data written over the life of the product, the system has to make sure that the entirety of the memory media is written more or less uniformly. This necessitates maintaining a logical-to-physical address mapping. Furthermore, it is likely that within a product, not all memory cells exhibit the same deterioration profile over time. Consequently, some areas of the memory will deteriorate sooner than others. Thus, health monitoring is necessary to maximize the potential of device usage.

3. Avoiding repeated read-related disturb. In some applications, the user application may request to read from the same memory word line location at disproportionately large amount of times. This may cause the NAND to provide a noisy output at the neighboring word lines. Suitable monitoring of read operation frequencies is then required for relocating data which it's neighboring WLs are frequently read appropriately.
4. Cross-temperature effects. In mobile products which may be exposed to a large scope of temperatures, the CVD may vary widely if cells were programmed at one temperature point and later on read at a different temperature. This phenomenon, as illustrated in Figure 2 above, occurs in 3D-NAND. To address this issue, the controller should monitor the temperature conditions during write for each physical region of the NAND, and apply proper compensation to read level thresholds when reading from a region programmed at a different temperature.

A generic structure for a flash-based system is shown in Figure 4 below. The scheme should address the challenges listed above.

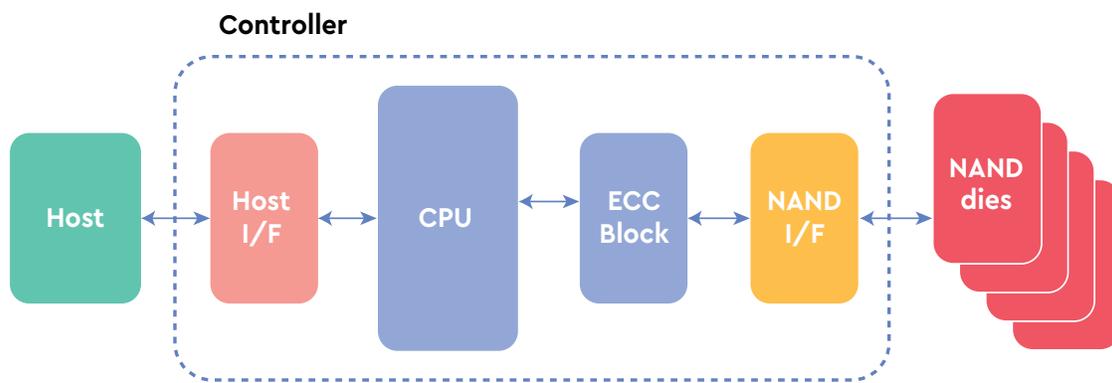


Figure 4: Generic structure for a flash-based system

In the next section of the white paper, you will find a high-level overview of our approach.

ECC-Centered Solutions

A strong ECC/DSP solution is typically required at the controller core as to address the challenges presented by a realistic NAND-based memory system. Generally, the ECC solution must have high performance while consuming low power, near-theoretical limit correction capability, and the ability to recover from failures. Depending on the application, a good quality of service, or latency profile, may also be required.

Western Digital's proprietary Sentinel ECC&DSP™ technology is embedded in all its NAND controllers. It is a mature technology with 15 generations deployed within Western Digital's controllers across various product lines (enterprise and client grade SSD, embedded NAND, memory cards, USB drives, etc.). A unique Sentinel ECC&DSP solution is tailored per product/application according to its specific requirements for throughput, latency, power, and other operational specs.

The Sentinel ECC&DSP error correction is based on state-of-the-art Low Density Parity Check (LDPC) coding and provides a full suite of NAND DSP (Digital Signal Processing) services, including data randomization or shaping, NAND health metering via Bit Error Rate (BER) estimation, ECC-based read thresholds calibration, and NAND defect protection and recovery via XOR based RAID scheme support.

The Sentinel ECC&DSP LDPC engine is generic and can support different memory types (SLC, MLC, TLC, QLC) in terms of the page size and ECC redundancy. The Sentinel ECC&DSP LDPC codes provide near Shannon limit correction capability and therefore maximize the supported operational specs (endurance, retention, disturb resilience, quality of service, etc.) for a given ECC overprovisioning in the NAND. The Sentinel ECC&DSP LDPC decoder is based on a proprietary multi-gear architecture, which optimizes the power consumption under variable BER observed across memory pages, operational conditions and memory lifetime. It supports multiple read

resolutions (for utilizing "soft" information). All transitions between the decoding gears are fully automatic based on internal BER estimation and require no firmware intervention. Overall, for a given set of product requirements, the Sentinel ECC&DSP engine has a very slim silicon area footprint and a low power usage.

Since during most of the memory lifetime the observed BER is low, the cost and power of the LDPC engine can be further optimized by employing multiple decoding gears with different power, throughput, and correction capability profiles.

Hence Western Digital's solution comprises of:

- Proprietary Bit-Flipping (BF) decoder, delivering high throughput with low power consumption with small silicon footprint.
- Additional decoding gears based on varying resolution fixed point Belief Propagation (BP) soft decoding.

An example of how said multi-gear engine performs is illustrated in Figure 5 below.

In this example, three decoding gears are used. The color-coding is used to show the BER regions where each gear is likely to succeed the decoding.

Gear 1 (the green region in Figure 5 below) is based on the bit-flipping decoder mentioned above. Said Gear-1 is characterized by high-energy efficiency (low J/GB/sec). However, as the BER goes up, it fails at some point to decode with high probability. At that point, the higher resolution decoding gears kicks in.

The blue curve in Figure 5 indicates the throughput of the system. Naturally, once NAND approaches the BER point where the 2nd gear kicks in, the throughput declines rapidly assuming constant power consumption is maintained.

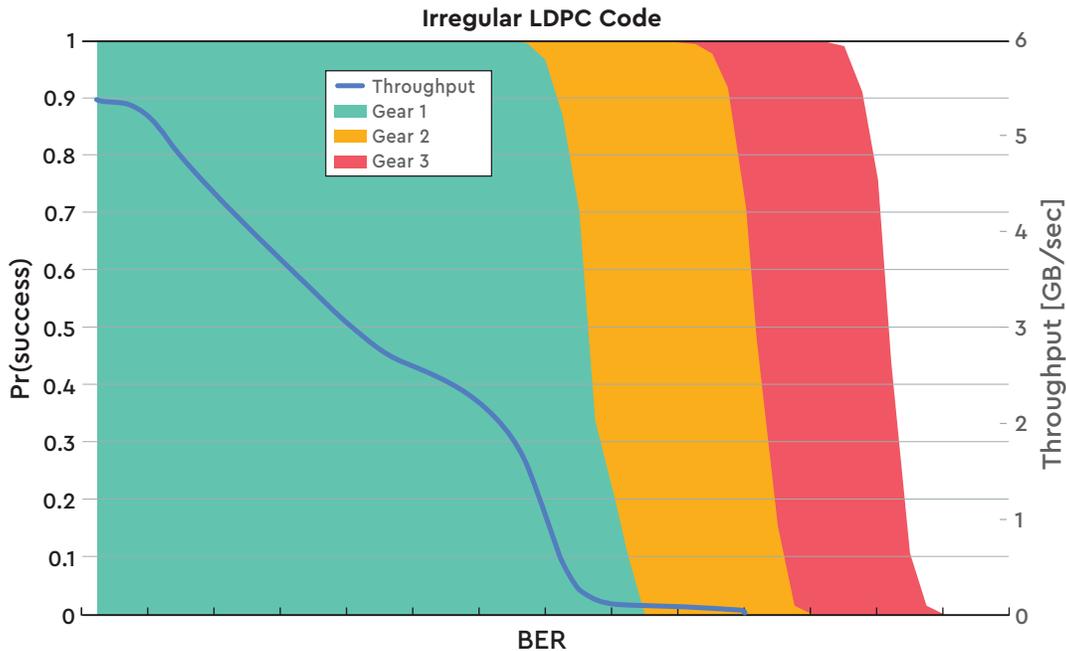


Figure 5: Example of various decoding gears correction capability and overall performance

When the BER is high, there is no point in spending decoder time trying to decode with lower gears, as this would only degrade performance by adding unnecessary latency to the decoding sequence. Consequently, the LDPC engine estimates the BER of the noisy page as part of its initialization process by counting the number of unsatisfied parity checks. This forms the basis on which it automatically chooses the appropriate decoding gear.

Furthermore, the parallelism of each decoding gear is dimensioned according to its usage probability (given the memory BER distribution) and the overall required decoding throughput. Thus, the number of costly high-resolution processing units instantiated for the full resolution BP decoder, which is rarely used ("safety net"), could be much lower than the number of simple BF processing units. This approach significantly reduces the ASIC footprint with a negligible impact on overall sustained decoding throughput.

Another important aspect, especially for enterprise grade solutions, is to maintain quality of service (QoS). This means that the delay of providing the host with response for any command falls under tight profile of probabilities it is allowed to exceed certain latency numbers. In order to meet such harsh requirements the ECC engine in gears 1,2,3 is segregated in a manner that it can serve separate requests in separate gears simultaneously. In case a single command requires higher latency gear-3 decoding, a different command can be attended simultaneously with Gear-1 engine. Together with out of order processing and host PCIe™ interface using NVMe protocol the Sentinel ECC&DSP technology is tailored to serve QoS in near optimal manner.

While the quality of the system solution depends heavily on the core competency of the ECC engine and its features, our experience has shown that other complementary DSP features embedded into the ECC engine can add tremendous value. Quite surprisingly, said DSP features – in aggregate – have negligible ASIC footprint.

RTC – Read Threshold Calibration

One of the main challenges for flash-based products from a systems point of view is the management and tracking of optimal read thresholds. Legacy methods primarily revolve around measuring the cell voltage distribution and thereafter estimating the minima points ("Valley Search"), between any 2 consecutive distributions selecting these as the optimal read thresholds. This method is illustrated in Figure 6 below.

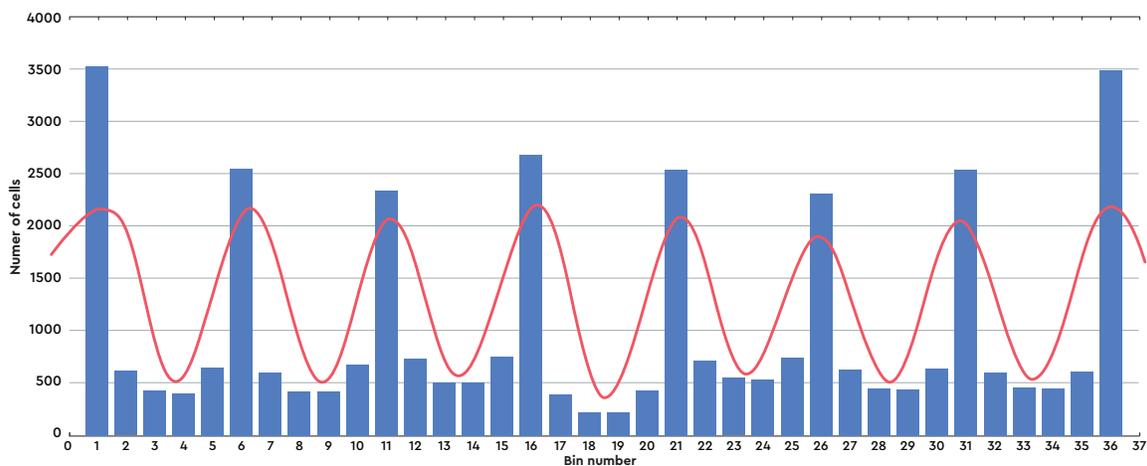


Figure 6: Cell voltage distribution estimate based on performing many reads

Western Digital's approach to this problem, as adopted in several product generations during the past years, is different.

Let us first note that Valley Search-based methods do not make use of the fact that the data in the flash is LDPC-encoded. Also, our experience show that these methods are not very robust when the variety of flash conditions and disturbs are considered.

Our RTC (Read Threshold Calibration) method (US patent #9,697,905) works as follows (Best explained using an example): Assume TLC flash for which the logical page we are interested in is read using thresholds between Er,A states plus thresholds between D,E states. We apply, in this example, 9 reads of this page with the same voltage window being shifted along values $[-4\Delta, -3\Delta, -2\Delta, -\Delta, 0, \Delta, 2\Delta, 3\Delta, 4\Delta]$. The results of these reads is stored in memory. Our objective is to find the optimal read thresholds among the $9^2=81$ possibilities. The key is to notice that the 9 reads, collectively, uniquely identify the voltage bin where each cell resides. This implies that we can digitally compute, for any combination of the 'left' e.g. threshold between 'Er' and 'A' states (9 possibilities) and 'right' e.g. threshold between 'D' and 'E' states (9 possibilities) thresholds, the resulting manifestation of the page data at any given combination of such potential selection of the two read thresholds.

Next, we rely on the fact that for each of these 81 possibilities, the data we are looking at is a noisy version of an LDPC code word. Thus, the syndrome weight (or number of unsatisfied check nodes) may be computed for all possibilities of choices of the read thresholds, with only 9 flash read operations (!). Now, since the syndrome weight (SW) is a proxy for the actual error rate, the optimal choice of read threshold coincides with the choice minimizing the SW. This process is depicted in Figure 7. Further digital processing of the result is also possible. Note that in Figure 7 only 5 read thresholds are used and not 9, the choice of how many such thresholds may depend on the expected range of the shift.

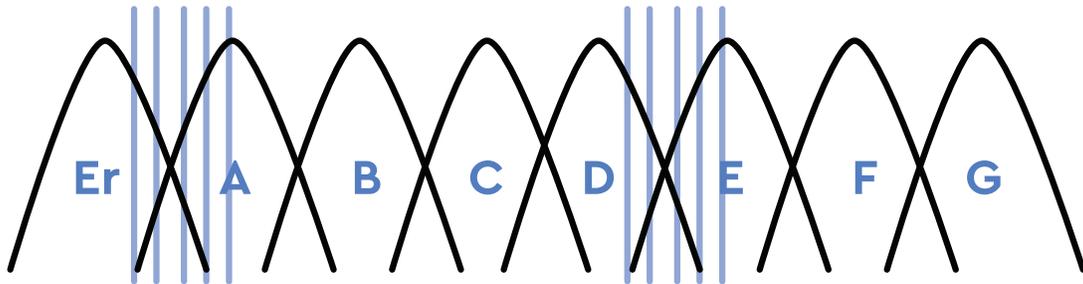


Figure 7: Example of RTC reads: TLC

The accuracy of RTC is shown in Figure 8 below. In this figure, RTC is compared against a legacy Valley Search method, with the optimal read thresholds also shown as a reference. In our experience, this example is representative of many cases, including extreme flash conditions. In other words, RTC has practically proven to be robust to varying flash conditions and is an indispensable tool used by the system for properly tuning flash read thresholds.

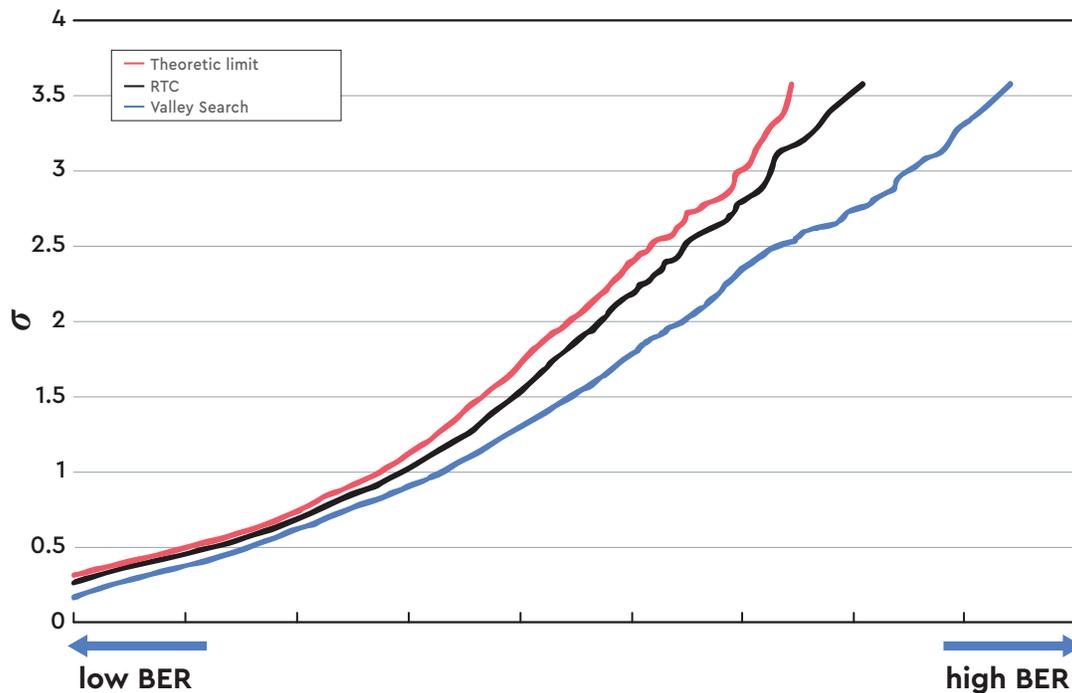


Figure 8: Comparison of RTC vs. Valley Search, optimal results shown as reference

Memory Error Model (MEM) Estimator

Standard message-passing decoding algorithms assume that the channel noise model is known perfectly by feeding the appropriate likelihood ratios or log-likelihood ratios (LLR) into the decoder. This assumes that a Gaussian noise distribution is a reasonable choice in most cases. However, under certain adverse flash conditions,

the voltage distribution can significantly deviate from a Gaussian distribution. This is illustrated in Figure 9. In such cases, where a stronger decoding gear must be used, potentially with soft bits, having at least an approximate model of the noise may be critical to decoding success. The effects of not knowing the channel model are not limited to a reduction in correction capability. Higher decoding latency is an additional possible outcome of decoding 'blindly'.

One possible approach to MEM estimation is to conduct offline estimation. This approach, however, requires maintaining a database of MEMs for different scenarios and then try them one by one. Said approach introduces significant overhead. A better approach is to adaptively estimate the error model for every specific memory page, as part of the decoding process. Since decoding of LDPC codes is iterative, the channel estimation mechanism can be incorporated into the iterative process, as shown in Figure 10 below.

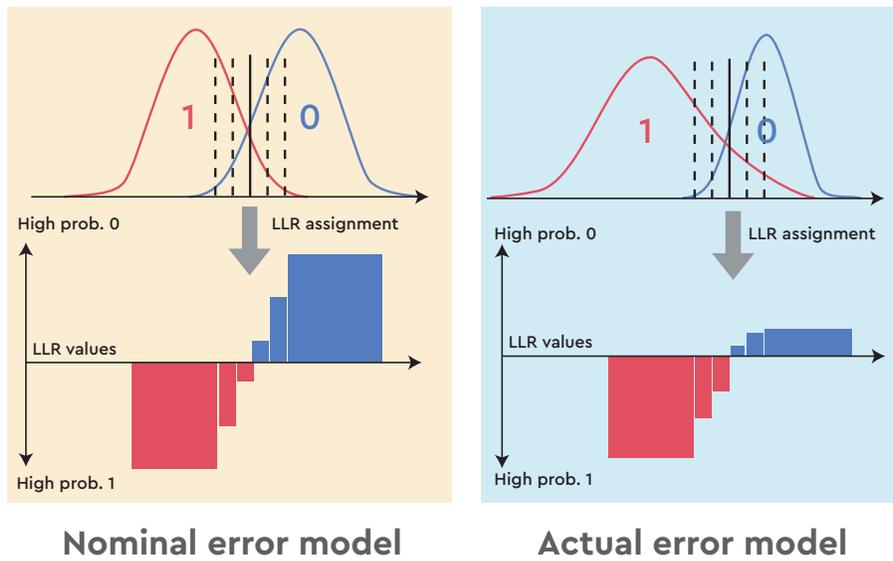


Figure 9: Channel error models

In Figure 9, we see on the left a nominal model, closely approximated by a Gaussian distribution. On the right is an asymmetric distribution not closely approximated as Gaussian.

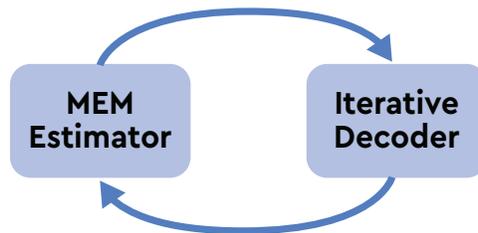


Figure 10: Iterations between decoder and MEM estimator

Denote the result of applying this iterative method the MEM Mismatched Decoder or MMD.

An example of the robustness of MMD to variance in the channel model is shown in Figure 11. On the left side, it shows the correction capability (black) plotted vs. the mutual information (MI). At each MI point, a Monte Carlo simulation consists of drawing a channel distribution for each code word at random, having that particular MI value and measuring the probability of decode failure. The channel distribution at each trial can be asymmetric. The yellow curve shows the theoretical limit e.g. providing the decoder optimal log likelihood ratio for each bit, while the blue curve shows the decoder correction capability when providing the decoder log likelihood thresholds representing gaussian channel and optimal read threshold.

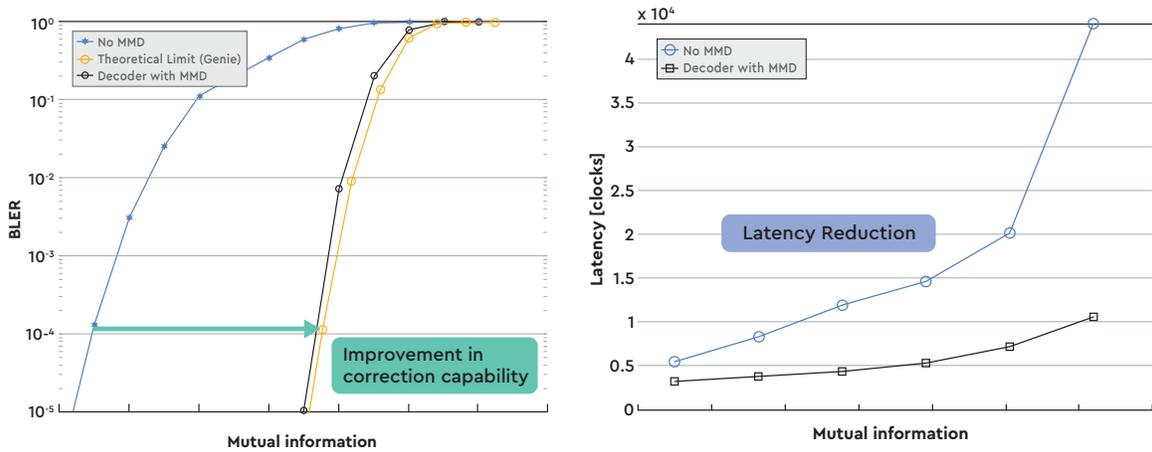


Figure 11: Performance under a random channel with fixed mutual information

Figure 11 above at the right hand side shows the latency under a random channel with fixed mutual information as a function of the mutual information, which increases towards the RHS (Right Hand Side) axis direction of the plot.

In Figure 11 on the left hand side for reference, a genie curve is included (yellow), where the correct decoder LLR values are known at the onset of decoding. Clearly, we cannot do better than this reference optimal curve. The blue curve shows the performance when no channel estimation is attempted. The fact that the MMD curve (in black) is very close to the genie curve is a strong indication of its robustness to variance in the channel model. Similarly, on the right side of Figure 11, we show the decoding latency under a same scenario. Clearly, significant latency reduction is achieved even when considering the extra latency introduced by the channel estimation algorithm itself.

We have earlier described our RTC algorithm, which is used to provide the ultimate read thresholds. But what if we need to read from a flash page where read thresholds are inaccurate, but the management algorithm at the system level has not yet corrected for this using RTC? Such a scenario demonstrates another benefit of the MMD scheme and illustrates how it nicely complements RTC-based read threshold management to provide great robustness.

When read thresholds are off from their optimal values, MMD can compensate for this, learn the skewed error model, and successfully decode while providing feedback to the system that RTC needs to be activated on this flash region. This is illustrated in Figure 12 below, showing that MMD can compensate when read thresholds are skewed away from their optimal value.

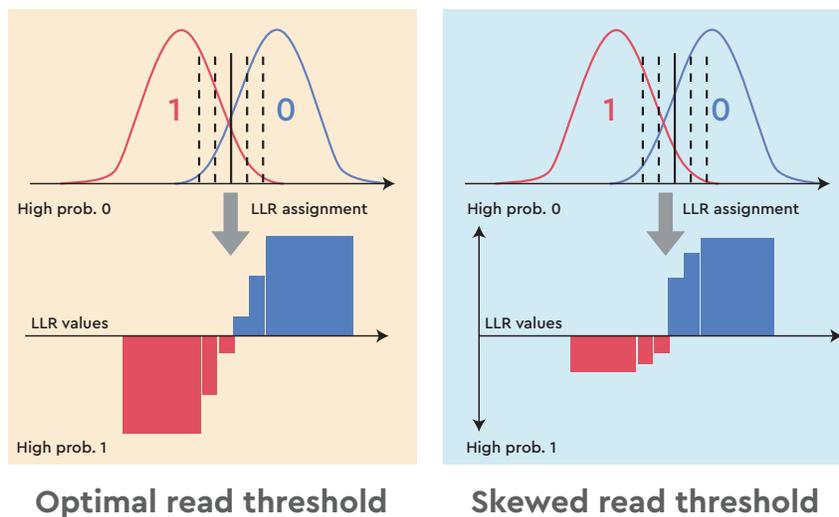


Figure 12: Read Thresholds

Joint LDPC and RAID Decoding for Enterprise SSD

Certain flash applications, such as enterprise SSD, require an exceptionally high level of data reliability. In these applications, the probability of losing data – no matter what caused it – must be exceptionally low. More specifically, the system must be able to cope with flash failure modes that do not only include random errors, but failures of entire pages. These might include word line failures of various sorts, block failures, or even die-level faults.

One existing strategy is to use RAID parity. This method has many variants depending on reliability requirements (how many failed sectors can the system recover from), locality of recovery (how many read and transfer operations required for a single recovery), and other system tradeoffs. The simplest of such schemes is the plain-vanilla variant, where recovery from a single sector failure is possible. This requires one parity sector, which holds the RAID parity of all other sectors. Upon failure of one of the sectors, reconstruction is possible by performing RAID recovery on all the non-failing sectors. Of course, this simple scheme is equally effective if the failing sector is due to a large number of random errors, which is not correctable by the ECC. The following two are simple observations on this scheme:

1. The ECC solution is completely decoupled from the RAID recovery layer.
2. The ECC cannot correct more than one failure, even if two failures are caused by a large number of random errors but not a complete fail.

To illustrate, please refer to the specific case study in the Appendix – Case Study at the end of this white paper starting on page 16.

Additional ECC/DSP features

Pseudo-Random Number Generation (PRNG). This enables generating short random sequences for services such as scrambling seed generation. This is not suitable for cryptographic applications.

Optional data shaping and scrambling. Data shaping allows for changing the distribution of states written to the NAND, which has the potential to improve flash endurance. This assumes that the incoming data is not encrypted.

BER estimation. Counting the number of unsatisfied check nodes will provide a proxy for the BER level. This proxy may be applied when the BER is beyond the decoding capability and may provide health metering service for various NAND management algorithms which need to make decisions based on the NAND health. These algorithms include wear leveling, deciding to relocate data based on highly frequent reads or other criteria, and others. BER estimation is also used as a building block for the RTC read threshold optimization scheme and automatic gear shifting mechanism.

RAID engine. A RAID store and recovery operation may be applied for protection against memory defects, supporting temporary snapshot generation, storage and reconstruction, or storing in a configurable manner either in temporary RAM or in SLC flash, depending on the application. It is a building block at the encode and decode phase for the combined LDPC-RAID scheme described above.

CRC. The standard error detection parity is appended to each written code word. It verifies post decoding in order to avoid undetected errors.

Puncturing and Shortening. As the ECC decoder is set to operate in high throughput the common layered approach is employed. This approach mandates to use layer granularity during encoding and decoding operations as to process simultaneously large number of bits. In order to maintain high resolution of information and parity shortening and puncturing operations are supported in both encoder and decoder modules. This feature is set to accommodate full usage of the available NAND cells in the physical page while enabling flexibility in code rate for numerous applications / scenarios.

Soft information translation. As described above the NAND media is characterized with large upper tail of FBC. In order to maintain system reliability soft information is read from the NAND, e.g. higher resolution read voltage bin is provided to the decoder which is translated to a reliability metric when ECC decode operation is activated. The Sentinel ECC&DSP solution is hence equipped with a translation interface module which converts the read threshold bin into log likelihood metric fed to both BP and BF gear shifts decoders. This feature not only increases the overall correction capability (can correct more bit flips per page) but can also be used to expedite decoding operation itself in order to reduce latency.

Summary

With each technology node progress, NAND flash is becoming smaller, denser, and more complex. This complexity breeds new challenges for storage systems based on flash, including higher error rates, higher failure rates, new failure modes, and increased sensitivity to operating conditions such as downtime, temperature, and endurance. This underscores the need for system-level solutions obligated to satisfy requirements across all product categories.

At the heart of the flash-based storage system is the ECC solution. The Western Digital Sentinel ECC&DSP LDPC solution complemented by a suite of DSP features, coupled with vertically integrated capabilities and technologies, provides an overall one-stop-shop package. This package can then be used by the system to tackle the many challenges posed by the NAND technology on the one hand and the product requirements on the other hand. This white paper doesn't cover the full system solution since that would be outside the scope. Our focus is on the DSP and ECC capabilities. Such a full system solution varies significantly between various product lines using different parts of the ECC/DSP package.

As flash technology continues to evolve, the ECC/DSP will adapt correspondingly to the storage technology and its future challenges.

One of the advantages of Western Digital as a lead player in NAND-based products is vertical integration. This capability enables many of the features described in this white paper. It is underpinned by the fact that the main technologies required to build a product are available internally within the company, including NAND design & manufacturing, controller ASIC design, and the FW/System design.

Appendix – Case Study [FMS 2016, Patent Application US2017025518A1]

Let's assume a standard RAID stripe of length 32. This means that 31 sectors are protected by an additional parity sector. The over-provisioning used by this RAID scheme is around 3%. Next, let's assume an LDPC-based ECC solution that has approximately 10% over-provisioning.

If the LDPC and RAID layers were to be combined, one could observe that the result is a code with extra 3% overprovisioning and a code word size that is x32 larger than the original. A Tanner graph representing a unified LDPC-RAID code is shown in Figure 13 below. It shows a Tanner graph of a combined LDPC-RAID scheme. The smaller graphs G_0, G_1, \dots, G_{30} represent the information sectors and G_{31} is the parity sector.

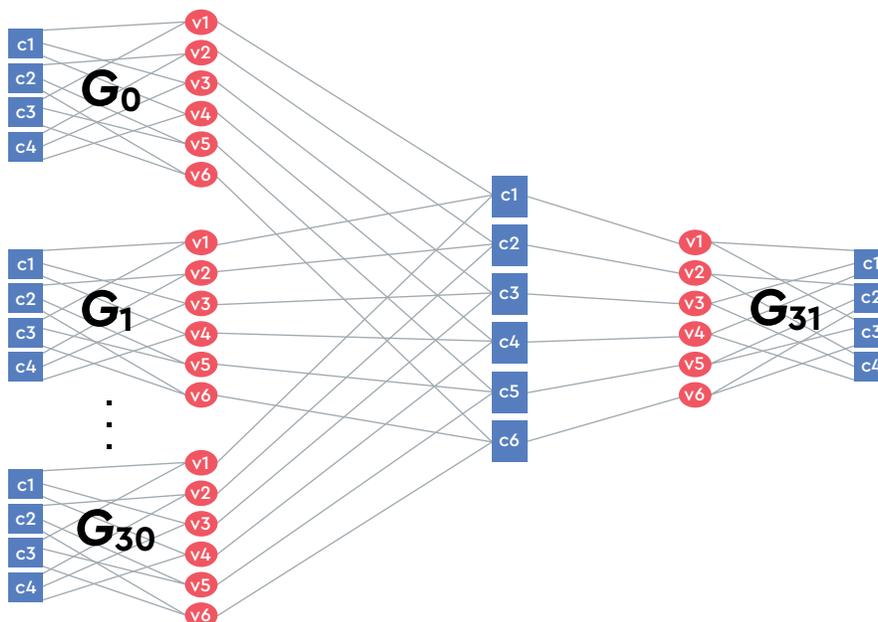
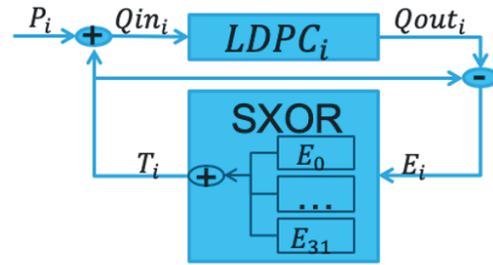


Figure 13: Tanner graph of a combined LDPC-RAID scheme

The graphical model of the combined LDPC-RAID scheme motivates the following recovery scheme, as shown in Figure 14.

In this recovery scheme, P_i is the vector of bit messages for sector i , which is to be reconstructed, expressed in the log-likelihood domain. The input and output bit LLR messages of sector i to the LDPC are denoted Q_{in_i} , Q_{out_i} , respectively. The extrinsic information for all code words $\{E_j\}_{j=0}^{31}$ is then combined to generate T_i (see Figure 14), which represents extra extrinsic information which can be added back to P_i , assisting to decode failing sector i .

This scheme can correct not only one decoding failure, but potentially up to 32 failures, by repeated iterations of the extrinsic information feedback loop of Figure 14.



$$T_i = \varphi^{-1}(\sum_{j \neq i} \varphi(E_j)),$$

$$\varphi(x) = \{\text{sign}(x), -\log \tanh \frac{x}{2}\}$$

Figure 14: Combined LDPC-RAID failure reconstruction scheme

A comparison of the LDPC-RAID scheme to a standard scheme is shown in Figure 15 below. As can be seen, the scheme discussed above improves considerably on the random-error correction capability of the decoupled LDPC-RAID scheme.

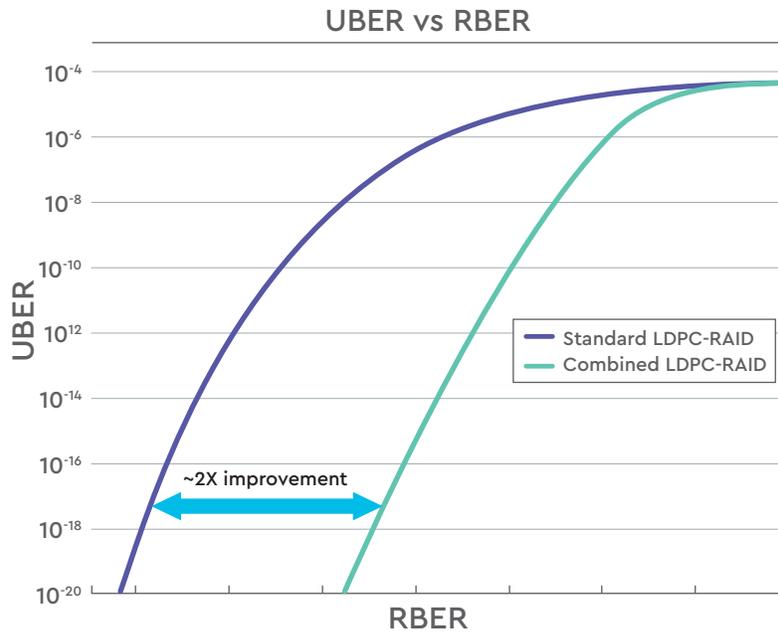


Figure 15: UBER curve for LDPC, optimal LDPC-XOR and simplified LDPC-XOR