

# Technical Disclosure Commons

---

Defensive Publications Series

---

December 2021

## Computing Optimal LDPC Decoding Voltages for Flash Memory

n/a

Follow this and additional works at: [https://www.tdcommons.org/dpubs\\_series](https://www.tdcommons.org/dpubs_series)

---

### Recommended Citation

n/a, "Computing Optimal LDPC Decoding Voltages for Flash Memory", Technical Disclosure Commons, (December 07, 2021)

[https://www.tdcommons.org/dpubs\\_series/4775](https://www.tdcommons.org/dpubs_series/4775)



This work is licensed under a [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/).

This Article is brought to you for free and open access by Technical Disclosure Commons. It has been accepted for inclusion in Defensive Publications Series by an authorized administrator of Technical Disclosure Commons.

## Computing Optimal LDPC Decoding Voltages for Flash Memory

### ABSTRACT


As flash memory cells age, the boundaries between cell voltages corresponding to ‘1’ and ‘0’ change. Measuring the movement with time of the voltage boundaries (or thresholds,  $V_{TH}$ ), and the confidence in the location of the boundaries (log-likelihood ratios, LLR) is critical to the reliable performance of the memory. This disclosure describes techniques for experimentally determining voltage thresholds and log-likelihood ratios values in flash memories, such that optimal  $V_{TH}$  and LLRs can be re-computed periodically or at predetermined ages. Once re-computed, the program-erase (P/E) cycle-optimized values can be used for reading and for soft decoding until the next re-computation of  $V_{TH}$  and LLR. The techniques can enable flash memories to have a longer effective lifespan and reliability.

### KEYWORDS

- Flash memory
- Solid-state drive (SSD)
- SSD lifespan
- Program/erase (P/E) cycle
- Log-likelihood ratio (LLR)
- Error-correcting code (ECC)
- Multi-level cell (MLC) flash
- Single-level cell (SLC) flash
- Long-term voltage drift
- Silicon aging
- Low-density parity check (LDPC)

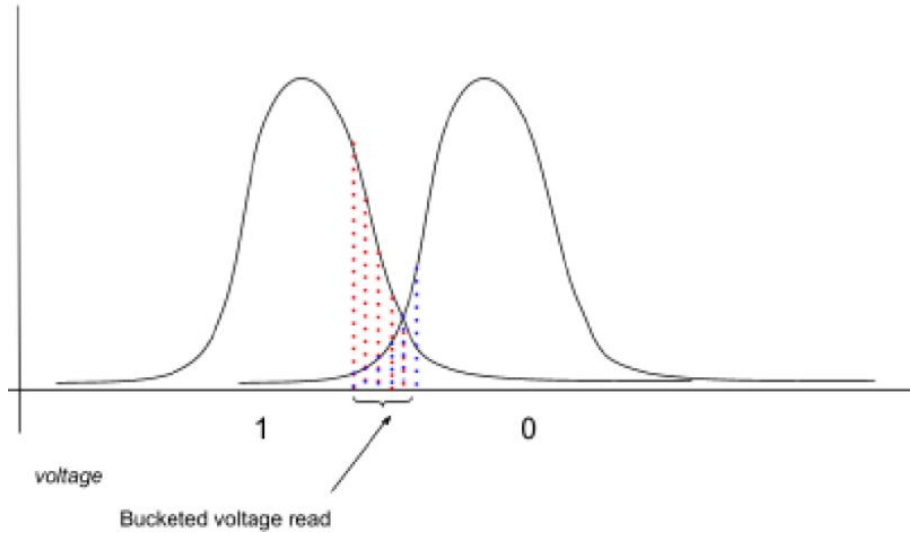
## BACKGROUND

Solid-state drives (SSD) that consist of flash memory are used in many applications that require fast read-write capability. Efficiently using SSDs entails the correction of data-read errors (or failures) caused by hardware, which usually worsen with the age of the SSD. Data is encoded using error-correcting codes (ECC) such as low-density parity check (LDPC) codes to protect it from errors or to enable recovery from errors during read-back. Optimal ECC performance is central to the effective use of SSD storage.



**Fig. 1: Voltage levels in a flash cell corresponding to different bit patterns**

Flash cells store data as voltages. In a simple case, certain voltages correspond to a '1' in memory, and others to a '0'. In a more sophisticated case, illustrated in Fig. 1, different voltage levels correspond to different stored bit patterns, as in a multi-level cell (MLC) flash. When a read operation is performed, a read-reference (threshold) voltage ( $V_{TH}$ ) is used to test the voltage present in the cell. A read operation is similar to comparing  $V_{TH}$  with the actual cell voltage. The comparisons are done until sufficient confidence is gained about the bit-pattern bucket within which the voltage likely lies. If the actual cell voltage is in a particular range, the log-likelihood ratio (LLR) is the log of the ratios of the likelihoods that the voltage is in a particular bucket.



**Fig. 2: LLR values given a bucketed voltage read**

For example, as shown in Fig. 2, the LLR value of the voltage (in the depicted bucket) being read as ‘1’ is given by

$$\text{LLR} = \log_2(\text{red area}) - \log_2(\text{blue area}). \quad (1)$$

Voltage buckets (boundaries) and LLR value buckets are typically provided by the flash memory vendor. However, as the cells age and log an increasing number of program/erase (P/E) cycles, the voltage distributions in the cells degrades, and the vendor-provided LLR values and  $V_{\text{TH}}$  buckets become suboptimal. Effectively, the boundaries between the voltages corresponding to ‘1’ and ‘0’ change with the age of the SSD.

Measuring the movement with time of the voltage boundaries (thresholds) is critical to the proper performance of the ECC. Estimating the LLR, which quantifies the confidence in the location of the voltage boundaries, is also critical to ECC performance. Whereas optimal  $V_{\text{TH}}$  and LLR values reduce read errors in a flash (which in turn reduces the load on upstream error-correction mechanisms such as RAID), optimal LLR tables and  $V_{\text{TH}}$  change over the life of a flash cell. Computing exact LLR values can be computationally challenging, and therefore

vendors provide only a few LLR tables, e.g., LLRs good for the beginning of life or for the end of life. Although  $V_{TH}$  and LLR for a particular SSD age can be estimated using the given vendor values and a mix of heuristics, such estimates are not based on mathematical principles and may not result in efficiency improvements.

## DESCRIPTION

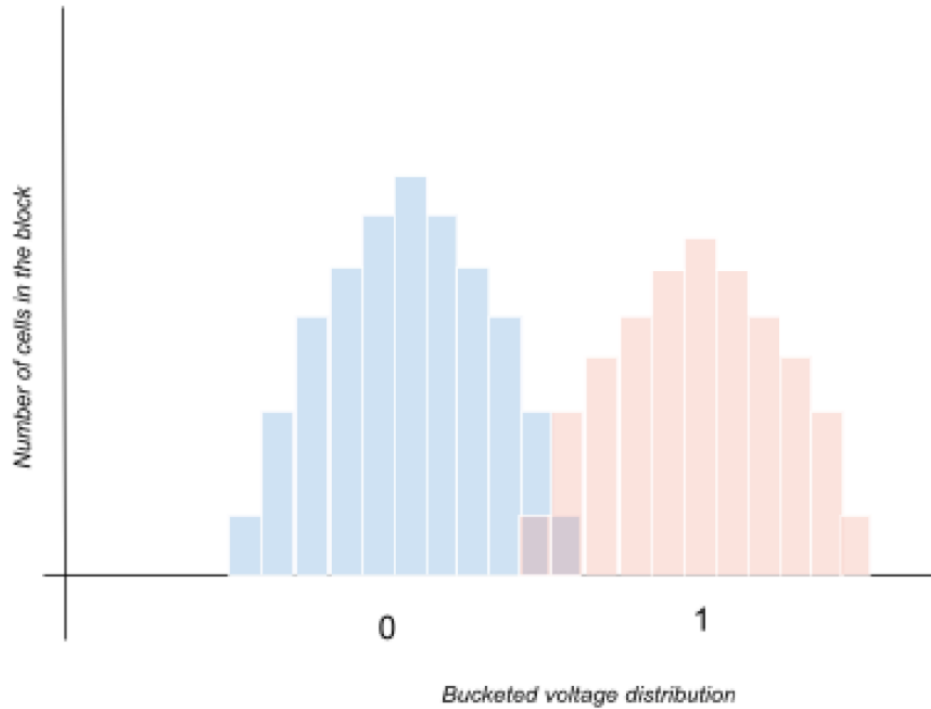
This disclosure describes techniques for experimentally determining  $V_{TH}$ s and LLR values in a flash memory. The techniques can be applied such that optimal  $V_{TH}$  and LLR values can be re-computed periodically, e.g., some multiple of one thousand P/E cycles, or at predetermined ages (lifecycle points). Once re-computed, the P/E-cycle optimized values are used for reading and for soft decoding until the next re-computation of  $V_{TH}$  and LLR. The experimentally optimized  $V_{TH}$  and LLR are referred to as instance optimal voltages and log-likelihood ratios, to clarify that these values are generated by experiment on a particular block, and are optimal (or nearly so) for that block. Blocks parameterized by similar P/E cycles and other relevant parameters can be expected to have similar instance optimal  $V_{TH}$  and LLR.

Experimental optimization of  $V_{TH}$  and LLR is cast as a statistical learning problem and divided into a data-collection phase and a distributional learning phase, as explained below.

### *Data collection*

A labeled example of the voltage distribution of a block is generated as follows.

- Write a predetermined string of bits into a block.
- Set read\_retry to the smallest step size, and read back all the data in the block.
- Create a histogram of the actual bucketed voltages labeled with the actual data it represents.



**Fig. 3: Example distributions (histograms) gathered experimentally. The blue histogram represents voltage values when a ‘0’ is written to the block; the red histogram represents voltage values when a ‘1’ is written to the block**

Fig. 3 illustrates example distributions (histograms) obtained from a single-level cell (SLC) flash. The blue histogram represents voltage values when a ‘0’ is written to the block; the red histogram represents voltage values when a ‘1’ is written to the block. The step-size of the histograms equal the smallest read\_retry, enabling the estimation of the actual voltages written into actual blocks. For this, and similar, blocks, the instance optimal LLR values and  $V_{TH}$  can be computed by solving simple learning problems on these labeled distributions, as described next.

#### *Distributional learning of $V_{TH}$ and LLR values*

Given a labeled, bucketed voltage distribution, the optimal  $V_{TH}$  and LLR values can be computed using standard techniques. For example, optimal  $V_{TH}$  can be computed using linear discriminant analysis (LDA) on the labeled data. The classification boundaries returned by the

LDA classifier are optimal if each distribution is Gaussian, a common assumption for the voltages in a flash cell. The  $V_{TH}$  thus obtained can be used for hard reads.

LLR values can be computed for each boundary as follows.

- For each bit pattern, compute the mean and the variance of its voltages. Model the bucketed distributions as Gaussian distributions parameterized by the computed means and variances.
- Knowing the means and variances of the Gaussian voltage distributions, the red/blue areas of Fig. 2, used to compute the LLR, can be computed using Monte-Carlo simulation, table look-up, or other standard techniques of Gaussian estimation. The LLR can be computed using Equation (1).

The LLR values derived from the above procedure quantify the confidence in the location of the voltage thresholds and can be used for soft reads.

Once the LLRs and the  $V_{TH}$  are found for an SSD of a particular age (lifecycle point), they can be used for other SSDs of the same age. For example, SSDs in a fleet from a given vendor can be classified by age, and LLRs and  $V_{TH}$  found for SSDs in each age class. The LLRs and  $V_{TH}$ , as a function of age, are applied to SSDs from that vendor across the fleet based on their age. Doing so can improve the performance of their ECCs and the reliability of the SSDs. The total number of iterations is the product of the number of vendors and the number of ages (lifecycle points).

Alternatively, the experimentally determined  $V_{TH}$  and LLR values can be used to validate the vendor-provided LLRs and  $V_{TH}$  values, or can be used to debug hardware issues in SSDs. Because the described techniques for determining  $V_{TH}$  and LLR rely on experiment rather than on relatively complex analytical computations, they are well-suited for practical implementation.

## CONCLUSION

This disclosure describes techniques for experimentally determining voltage thresholds and log-likelihood ratios values in flash memories, such that optimal  $V_{TH}$  and LLRs can be re-computed periodically or at predetermined ages. Once re-computed, the P/E-cycle-optimized values can be used for reading and for soft decoding until the next re-computation of  $V_{TH}$  and LLR. The techniques can enable flash memories to have a longer effective lifespan and reliability.