

Characterization and Analysis of Bit Errors in 3D TLC NAND Flash Memory

Nikolaos Papandreou*, Haralampos Pozidis*, Thomas Parnell*, Nikolas Ioannou*, Roman Pletka*, Sasa Tomic*, Patrick Breen†, Gary Tressler†, Aaron Fry‡, Timothy Fisher‡

*IBM Research - Zurich, Säumerstrasse 4, 8803 Rüschlikon, Switzerland

†IBM Systems, 2455 South Road, Poughkeepsie, NY, USA

‡IBM Systems, 10777 Westheimer Road, Houston, TX, USA

Abstract—3D NAND flash memory has entered dynamically into the space of enterprise server and storage systems, offering significantly higher capacity and better endurance than the latest 2D technology node. Moreover, the advancements in vertical stacking, cell design and program/read algorithms, have also enabled TLC 3D NAND flash with enterprise-level reliability, thus achieving further increase in capacity and cost-per-bit reduction. This paper presents an in-depth analysis of the bit-error characteristics of state-of-the-art 64-layer 3D TLC NAND flash with a focus on read-voltage calibration. We provide experimental measurements of the RBER and threshold voltage distributions using typical and mixed-mode test patterns of program/erase cycling, retention and read-disturb. Moreover, we quantify the RBER components attributed to threshold voltage level overlapping and on-chip 2-step program errors. Finally, we characterize how the optimal read voltages change under different device stress and we evaluate calibration schemes with different performance and complexity trade-offs.

Index Terms—3D NAND, flash memory, triple-level cell (TLC), endurance, threshold voltage.

I. INTRODUCTION

Flash memory technology has brought a revolution in consumer electronics, offering vast data storage capacity at small footprint, thus enabling miniaturization and numerous new applications. More recently, flash also made inroads into the enterprise storage space and the datacenter, creating a new tier between fast, volatile main memory and slow, nonvolatile hard disk drives. The inexorable demand for more storage capacity on one hand, and the continuous semiconductor technology node scaling on the other have fueled a sustained improvement of flash device capacity over the last years. The latest episode is the industry transition from 2D to 3D NAND. In addition to a capacity boost, the inherent relaxation of the cell lithography in 3D flash, brought a significant improvement in reliability, most notably endurance. Further advancements in vertical etching, layer processing and program algorithms have enabled triple-level cell (TLC) 3D NAND flash with enterprise-level reliability, thus further boosting capacity and reducing cost-per-bit [1]–[3].

Various studies, mainly on 3D MLC NAND, but also on 3D TLC NAND devices, have reported endurance and retention measurements, and also discussed differences between 2D and 3D NAND reliability aspects [4]–[8]. In this paper, we present an in-depth analysis of the bit-error characteristics of state-of-the-art 64-layer 3D TLC NAND flash. First, we provide

experimental measurements of the raw bit-error rate (RBER). We examine the distribution of errors across the different pages and page types of the NAND block using typical and mixed-mode test patterns of program/erase (P/E) cycling, retention and read-disturb. Moreover, we quantify the RBER components attributed to errors resulting from the overlapping of adjacent threshold voltage (V_{TH}) levels or originating from the inherent on-chip 2-step program process. Moreover, we characterize the V_{TH} distributions and analyze how the optimal read voltages change under different device stress. We further discuss the variability of optimal read voltage settings as a function of the page location in the block. Finally, we evaluate read-voltage calibration schemes with different performance and complexity trade-offs.

The presented analysis provides useful insights regarding the degradation characteristics of 3D TLC NAND flash under various stress conditions and can facilitate the development of effective error mitigation and fault-tolerance techniques. The benefits of such a comprehensive bit-error analysis as the one presented in this paper span various aspects in flash controller development, in particular read-voltage calibration, error mitigation and data management, as will become clear in the following sections.

II. 3D TLC NAND BLOCK AND PAGE PROGRAM ORDER

A block of 3D NAND flash comprises a layered structure of interconnected memory cells. Let us consider a block with L layers, in which each layer consists of a grid of N_b by M cells where N_b is the number of bit-lines (BLs) and M is the number of word-lines (WLs) per layer. Such an architecture is illustrated in Fig. 1 for an exemplary block with 8 layers and 8 WLs per layer. In 3D TLC devices, each cell stores three bits of information, which we will refer to henceforth as B_0 , B_1 and B_2 , where B_0 is the least-significant bit (LSB) and B_2 is the most-significant bit (MSB). Data are programmed to and read from the block in units of pages, which comprise a single bit from each cell in said WL. The length of each page is thus N_b bits and every WL contains exactly three pages. The total number of pages per block is therefore given by $3LM$, which for the example in Fig. 1 corresponds to 192 pages, but in real devices is typically more than a thousand.

A page is programmed by iteratively applying a sequence of programming pulses that alter the threshold voltage of




Fig. 1. Exemplary 3D TLC NAND block with 8 layers and 8 WLs per layer.

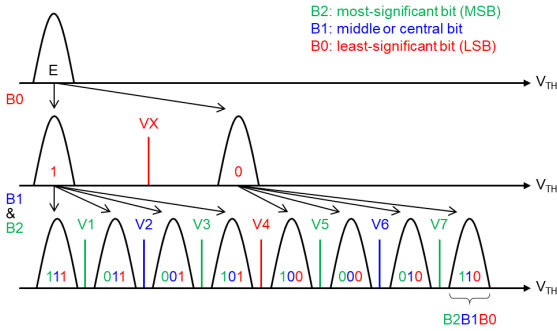

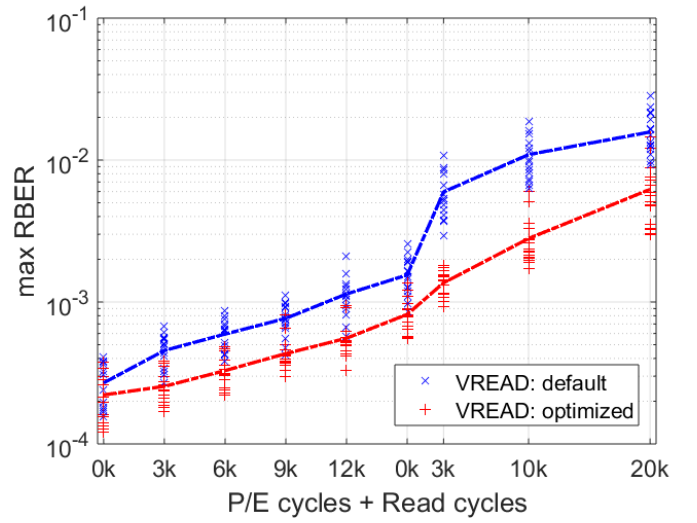


Fig. 2. Example of 2-step program and 3-bit readout in 3D TLC NAND.

the cells in the corresponding WL. It has been shown that, due to cell-to-cell interference (CCI) effects, programming a page from WL_i in layer L_k , can adversely affect the threshold voltage of the cells in WL_i in neighboring layers L_{k-1} and L_{k+1} . To mitigate this effect, device manufacturers specify a carefully-engineered page programming order. The exact order varies depending on the manufacturer, but one such order is illustrated in Fig. 1. In this order, after programming page B0 from layer L_k the NAND controller proceeds to program page B0 from layer L_{k+1} , before returning to layer L_k to program pages B1 and B2. The benefit of such a scheme is that the multi-step procedure for programming the higher-order bits is able to account for the CCI that takes place when programming pages in adjacent layers. We note that in such a scheme, in order to correctly program the higher-order pages, the device must establish what data was previously written to B0 page. This can either be achieved by the user providing the cached pages themselves, or the B0 data can be read internally by the device during the B1 and B2 program phase. This process is illustrated in Fig. 2, whereby the device internally applies a read voltage (V_X) to determine what data were written to B0, before proceeding to program the cells to one of 8 target threshold voltage levels. Once a WL has been finalized, i.e., all three pages have been programmed, the pages can be read



(a)



(b)

Fig. 3. Multi-block measurements showing the maximum page-RBER with default and optimized read voltages: (a) 9k P/E cycles followed by 3 months retention, (b) 12k P/E cycles followed by 20k block read-only cycles.

by applying a set of 7 read voltages, as illustrated in Fig. 2. In this example, page B0 can be read by applying a single read voltage (V_4), page B1 requires two read voltages (V_2 , V_6), and to retrieve information from page B2 requires four read voltages (V_1 , V_3 , V_5 , V_7). 3D TLC manufacturers provide commands that can be used to apply different offsets to the default read voltages, allowing the NAND controller to adapt to changes in the V_{TH} distributions that occur over time or under certain operating conditions.

III. RAW BIT-ERROR RATE CHARACTERISTICS

Fig. 3 shows the RBER when the device is subjected to P/E cycles followed by data retention idle time or read-disturb cycles. For each case, we randomly selected 24 blocks corresponding to different planes and different packages. In all cases, testing is performed with a cycling acceleration factor

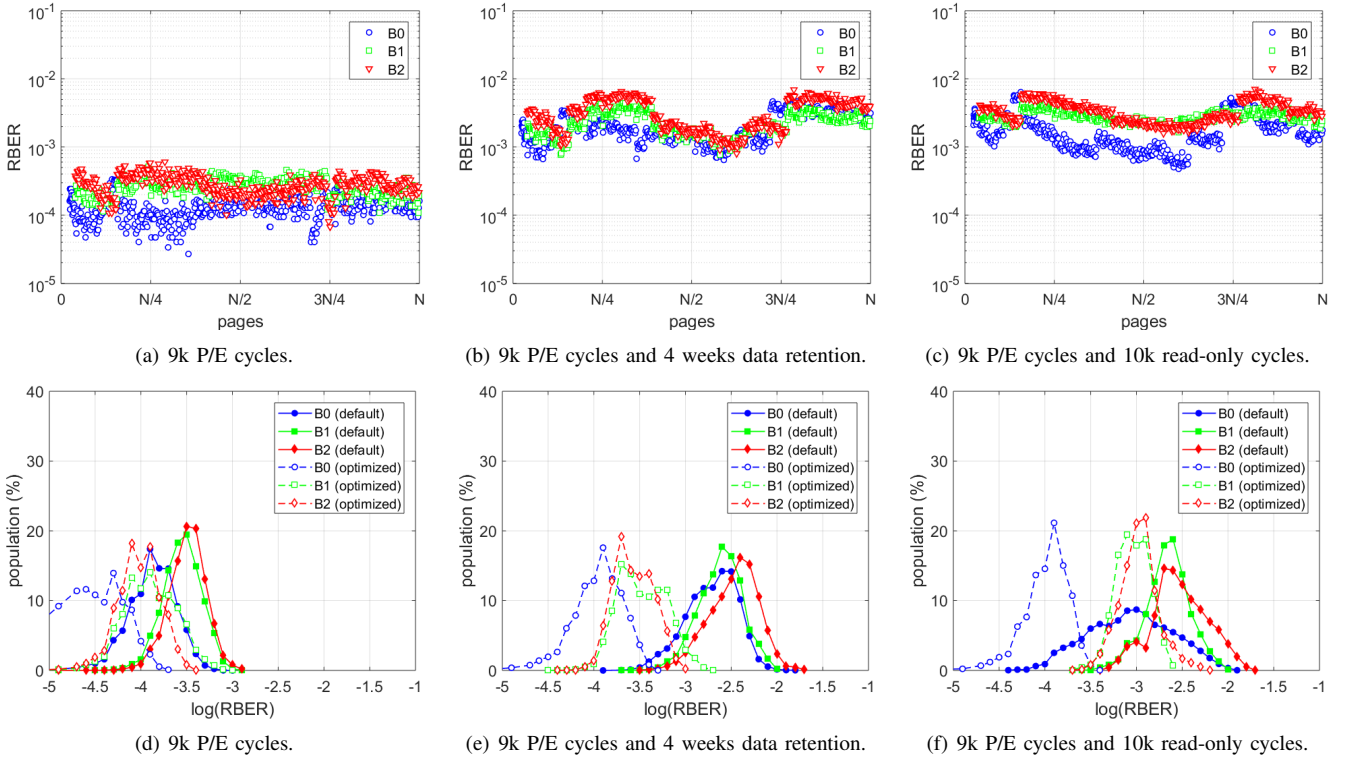


Fig. 4. RBER characteristics per TLC page type (B0, B1, B2) for different device stress conditions. [Top row] RBER as a function of the TLC page address in one 3D NAND block under test. [Bottom row] Histogram of the page-RBER (in log) of all blocks under test using default and optimized read voltages.

of 80/20, which implies that the first 80% of the targeted P/E cycles are tested using a minimal dwell-time, whereas the last 20% are tested using the nominal value [9]. In our test, we applied the 80/20 rule in every cycling interval between measurement collection, i.e., every 3k P/E cycles. The testing sequence runs under elevated temperature to accelerate data retention and charge-loss recovery phenomena, where we used the activation energy provided by the NAND chip supplier.

In Fig. 3(a), the blocks were first subjected to 9k P/E cycles with a nominal dwell-time of about 1600 sec and then the RBER was monitored at regular intervals corresponding to a total equivalent time of 3 months at 40°C. In Fig. 3(b), the blocks were first subjected to 12k P/E cycles and then read-only cycles were applied, where in each cycle we read all pages in the block. In this paper, read-only cycles will adhere to the previous definition. For all cases, we collected data using a read-voltage sweep. This enables the calculation of optimal offset values that return the minimum RBER under the different stress conditions. In addition, it allows the extraction of the V_{TH} distributions as will be shown in Section IV. Fig. 3(a) and Fig. 3(b) show the maximum page-RBER in each block under test using the default (in blue) and optimized read voltages (in red). The average value across all tested blocks is also shown (dashed line). We observe that the RBER increases gradually during P/E cycling and more abruptly in the data retention and read-disturb phases. The much steeper increase of RBER during data retention in 3D NAND was also reported in [4]–[8]. Note that the target temperature is different between

cycling and retention/read-disturb. This leads to an inherent delay for proper temperature regulation that explains the slight increase of the RBER at the end of the cycling period and the beginning of the subsequent data retention or read-disturb phases.

From the results in Fig. 3, we observe that the RBER increase is much smaller and less abrupt with optimized read voltages, in particular during data retention where the biggest gains from using the optimal offsets are observed. On the other hand, less gains are achieved in the high read-disturb regime, i.e., above 10k read-only cycles. This behavior is explained by the analysis of the extracted V_{TH} distributions in Section IV. Furthermore, we observe that the variability between blocks is also reduced with optimized read voltages. Here, we calibrated the read voltages for each page separately. Due to the large number of pages (the size of the block in state-of-the-art 3D TLC NAND devices is in the order of 1k-2k pages), it may not be feasible to keep meta-data of optimized voltage values for each page in a block. In this case, grouping of pages of the same type and with similar bit-error and read-voltage characteristics, where the same offset values are applied to all pages in a group, can be beneficial. Such a strategy relaxes the amount of meta-data, however it may result in an increase of the maximum RBER depending on the uniformity of the page characteristics in each group.

The bit-error characteristics of the various TLC page types (B0, B1, B2) are examined in Fig. 4 for three selected readouts of Fig. 3: (a) 9k P/E cycles, (b) 9k P/E cycles and 4 weeks

data retention, (c) 9k P/E cycles and 10k read-only cycles. The measurements in the graphs of the top row are from three different blocks using the default read voltages. In Fig. 4(a), we observe that immediately after programming at 9k P/E cycles, all pages exhibit fairly uniform bit-error characteristics with the B0 pages having a lower RBER level. Fig. 4(b) shows that the RBER of all pages increases significantly to a high value close to 10^{-2} as the result of charge loss after 4 weeks data retention. We also observe that the various pages types (B0, B1, B2) exhibit a similar RBER profile, however, the bit-error characteristics per page vary depending on the page index in the block. The first observation is explained in the next Section by examining the effect of charge loss on the programmed V_{TH} distributions. The second observation is related to process variations between different layers and, in general, between different locations in the block. This behavior suggests that read-voltage calibration, where the same offset values are used for a subset of pages, should be based on proper location-aware grouping of pages. Such an information can be collected in advance by large-scale lab measurements or can be computed and adapted on-the-fly during device operation. Fig. 4(c) shows the RBER measurements per page after 9k P/E cycles followed by 10k read-only cycles. We observe that the overall RBER has increased and the various pages exhibit non-uniform RBER characteristics as well. The effect of read-disturb mechanism on the different V_{TH} levels is discussed in details in the next Section.

As we have seen in Fig. 3, proper adjustment of the read voltages is necessary to reduce the maximum RBER to a level below 10^{-2} . In this regime, the bit-error rate can be tolerated with state-of-the-art error-correcting codes (ECC) that are typically used in storage systems [10]. In the bottom row of Fig. 4, we show the page-RBER histogram with default and optimized read voltages using data from all the blocks in each test. Fig. 4(d) shows that at 9k P/E cycles the overall RBER is below 10^{-3} using the default read voltages and can be reduced further with read-voltage calibration. In addition, we observe that the RBER histograms of the B1 and B2 pages are quite similar, whereas the B0 pages have relative lower error characteristics. In Fig. 4(e), we observe that after 4 weeks data retention the RBER is increased to high values close to 10^{-2} and all three page types (B0, B1, B2) have similar RBER histograms when the default voltages are used. The RBER characteristics improve significantly with optimized read voltages. Fig. 4(f) shows that after 10k read-only cycles the RBER increases to a high level close to 10^{-2} as in Fig. 4(e), however the RBER profile of the TLC pages (B0, B1, B2) is different. More specifically, the B0 pages show a large spread when the default read voltages are used. With optimized read voltages, there is a substantial improvement of the B0 pages and a relative less improvement of the B1 and B2 pages. These observations provide useful insights on how the different TLC page types are affected from the various stress conditions and how read-voltage calibration is able to reduce both the maximum RBER as well as the RBER variability between page types and page indexes in a block.

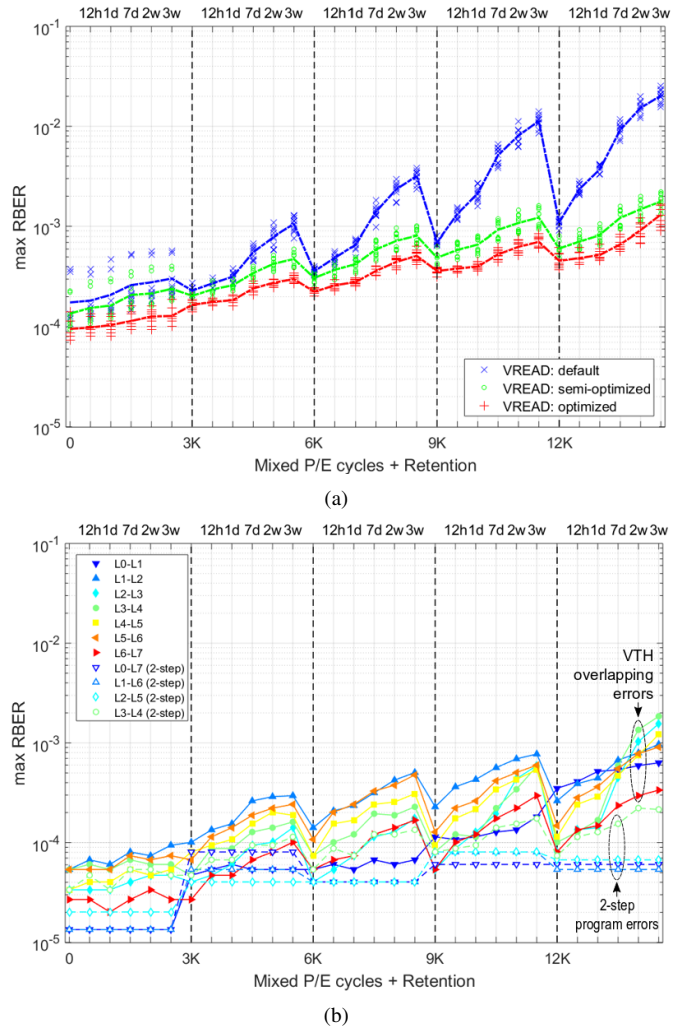


Fig. 5. Multi-block measurements during successive intervals of mixed 3k P/E cycles and 3 weeks data retention: (a) comparison of different read voltage optimization approaches; (b) comparison of level-to-level bit-errors using optimized read voltages (data from one block).

IV. ANALYSIS OF THRESHOLD VOLTAGE DISTRIBUTIONS

In this section we analyze the cell threshold voltage distributions. In order to collect valuable data that allow the extraction of the V_{TH} distributions and the calculation of the optimal read voltages as the device undergoes different types of stress, we designed a mixed cycling and retention experiment as follows: we randomly selected 12 blocks and applied successive intervals of 3k P/E cycles with 3 weeks data retention for a total of 12k P/E cycles. In contrast to the typical testing procedure of P/E cycling followed by data retention that was presented in Section III, the particular mixed cycling and data retention experiment provides useful insights on how the V_{TH} distributions and the optimal read voltages change during successive phases of different stress type. This information can guide the development of read-voltage calibration algorithms that are able to track the changes of the V_{TH} distributions and adapt the read voltages accordingly, so that the overall RBER is minimized.

Fig. 5(a) shows how the maximum RBER of each block changes during the successive cycling and retention phases. We observe that the RBER increases significantly during data retention and then drops during P/E cycling to a lower level that reflects the permanent wear of the block as a result of the P/E cycling history. The average value among all blocks is also shown (dashed line). The graph compares the maximum RBER using the default (in blue), optimized (in red) and semi-optimized read voltages (in green). The latter is a suboptimal scheme with a single corrective offset per page type. Such a scheme can reduce the required metadata, however, as the same offset is applied to all read voltages of a page type, an increase in the maximum RBER is observed. Moreover, care must be taken when the device stress conditions have a different effect on each V_{TH} level and thus a single offset may be not effective. In this case, the controller may need to resort to more frequent calibrations or data movement operations. Fig. 5(b) shows the RBER decomposition into bit-errors attributed to the overlapping of adjacent V_{TH} levels and bit-errors that originate from the on-chip 2-step program process described in Fig. 2. The latter type of errors is visualized in the measured V_{TH} distributions.

Fig. 6 compares the 8 discrete V_{TH} levels (L0-L7) at selected points of the mixed cycling and data retention test in Fig. 5 and of a similar mixed cycling and read-disturb test with successive intervals of 3k P/E cycles followed by 6k read-only cycles. Each graph shows the V_{TH} distributions within a window of threshold voltage values that is constrained by the range of the offset values used in the read-sweep commands. As a result, the graphs illustrate only the right tail of level L0 (erase), and the left body of level L7. In all histograms, the solid vertical lines mark the default read voltages and the threshold voltage axis is shown in arbitrary units as the actual values are not available.

Fig. 6(a) shows that after P/E cycling the distributions become wider and L0 develops a longer right tail. The widening of the distributions results in increased overlap between the adjacent levels, which in turns lead to an increase of the RBER. We also observe the presence of 2-step program errors that result in displacement of cells to different V_{TH} values, e.g., L0-to-L7. These type of errors originate from LSB mis-detection during the on-chip 2-step process when programming the higher-order bits [11], [12]. Fig. 6(b) shows that during data retention all distributions exhibit a negative shift as a result of charge loss and a further broadening. We observe that the default read voltages are no longer appropriate and a negative corrective offset that takes into account the effect of charge loss needs to be applied. Fig. 6(c) shows that after the application of read-disturb cycles the lower V_{TH} levels, i.e., L0 and L1, develop longer right tails. This behavior is explained by the fact that read-disturb results in over-programming due to the application of high voltages to de-select all other pages except from the target one. These high voltages mainly affect the cells that have been programmed to a low V_{TH} level. A moderate negative shift attributed to short-term retention is also present in all levels.

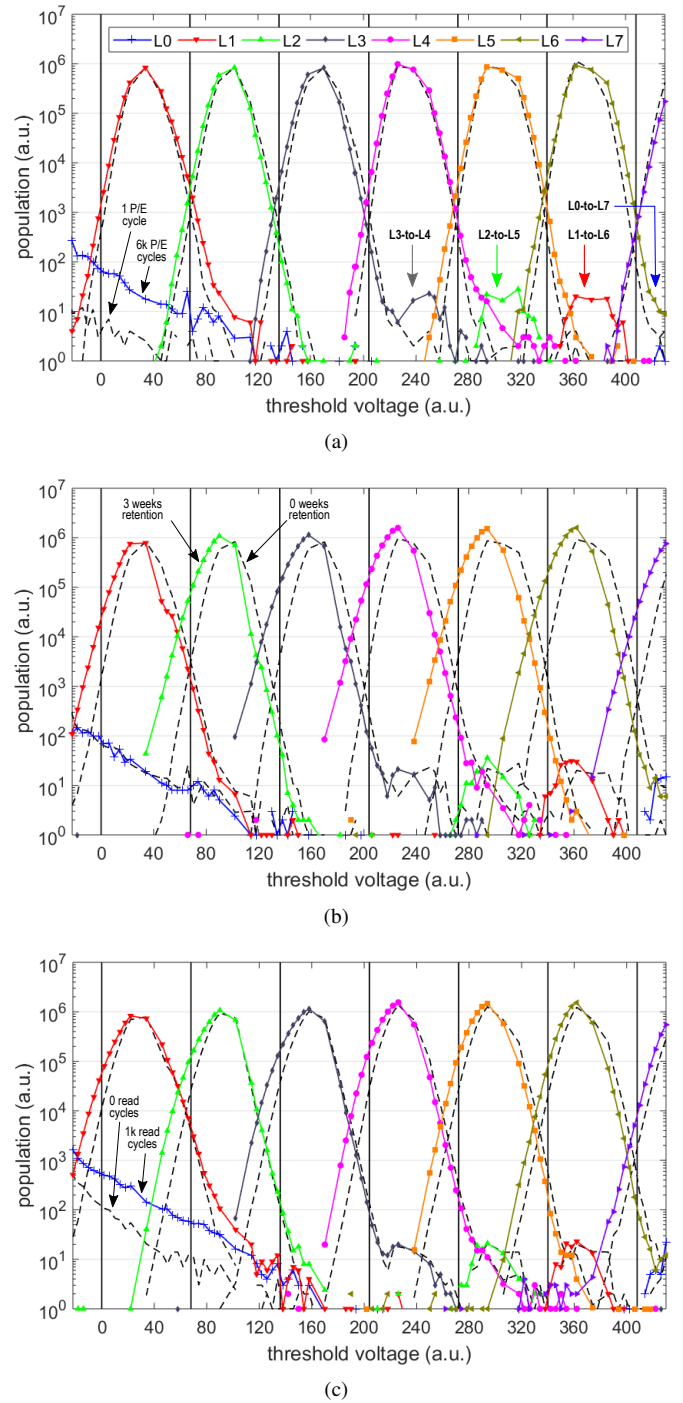


Fig. 6. Measured V_{TH} distributions: (a) 1 vs. 6k P/E cycles; (b) 6k P/E cycles with 0 vs. 3 weeks retention; (c) 9k P/E cycles with 0 vs. 1k read-only cycles.

The previous graphs highlight the fact that read-voltage calibration is crucial for keeping a low RBER and thus maintain the reliability of the device. Moreover, the presented analysis shows that the different types of stress have a different effect on the various V_{TH} levels. Further, it is shown that the 2-pass program errors cannot be mitigated by read-voltage calibration, however their contribution to the overall RBER was small as it is illustrated in Fig. 5(b).

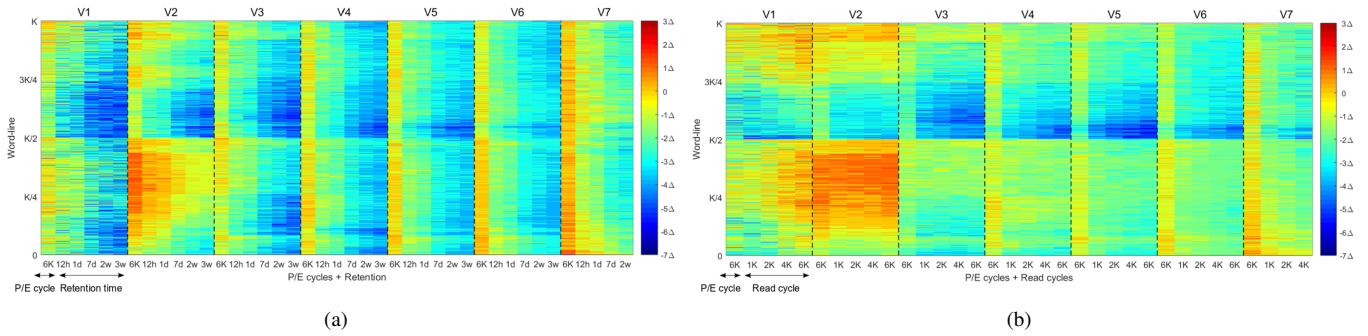


Fig. 7. Heatmaps of optimal read-voltage offsets showing the different requirements in optimization and tracking between the various read voltages (V1-V7) for different device stress conditions: (a) 6k P/E cycles followed by a 3 week data retention interval; (b) 6k P/E cycles followed by a 6k read-only cycles interval. The heat color reflects the positive (red) and negative (blue) offsets used in the read-sweep commands.

V. ANALYSIS OF OPTIMAL READ VOLTAGES

In this section, we analyze the optimal read voltages and examine the variability and dependency of the optimal offset values on the page index and page location in the block. Fig. 7 shows how much and towards which direction the optimal read voltage offsets change between successive intervals of P/E cycles and retention (Fig. 7(a)) or read-disturb (Fig. 7(b)). We present a single readout point at 6k P/E cycles from the corresponding mixed P/E cycling with retention and read-disturb experiments described in the previous Section.

The reported heatmaps illustrate the variability of the optimal offset values depending on the WL location in the block. These characteristics are attributed to various factors including, but not limited to, the trimming of the default reference voltages that may be different for the different areas in the block, as well as the effect of process variations that leads to variability on the programmed threshold voltage distributions in the various WLs, which in turns results in different values for the optimal corrective offsets. Moreover, we observe that the amount and range of the optimal offset values is different between the seven read voltages V1-V7, which explains the difference in RBER performance between the optimized and semi-optimized calibration schemes presented in Fig. 5(a). Finally, both graphs demonstrate the different requirements in optimization and tracking for the various read-voltages (V1-V7), between the two different types of device stress. This result was explained by the different effect of data retention and read-disturb mechanisms to the various threshold voltage levels L0-L7 that was discussed in Section IV.

VI. CONCLUSION

In this paper, we presented an in-depth analysis of the bit-error characteristics of state-of-the art 64-layer 3D TLC NAND flash. We provided experimental measurements of the RBER and of the threshold voltage distributions using typical and mixed-mode test patterns of program/erase cycling with retention or read-disturb. Moreover, we characterized how the optimal read voltages change under the various device stress conditions and we quantified the component of bit-errors originating from the overlap between adjacent levels and from

the on-chip 2-step program process. Finally, we evaluated both optimized and sub-optimized read-voltage calibration schemes and we discussed their different performance and complexity trade-off based on the analysis of the measured threshold voltage distributions. The presented study provides useful insights of the bit-error characteristics of 3D TLC NAND and can guide the development of effective read voltage calibration and error mitigation techniques.

REFERENCES

- [1] H. Kim, S. Ahn, Y. G. Shin, K. Lee, and E. Jung, "Evolution of NAND Flash Memory: From 2D to 3D as a Storage Market Leader," in *2017 IEEE International Memory Workshop (IMW)*, pp. 1-4.
- [2] S. Inaba, "3D Flash Memory for Data-Intensive Applications," in *2018 IEEE International Memory Workshop (IMW)*, pp. 1-4.
- [3] K. Parat and A. Goda, "Scaling Trends in NAND Flash," in *2018 IEEE International Electron Devices Meeting (IEDM)*, pp. 2.1.1-2.1.4.
- [4] P. Breen, T. Griffin, N. Papandreou, T. Parnell, and G. Tressler, "3D NAND Assessment for Next Generation Flash Applications," in *2016 Flash Memory Summit*.
- [5] K. Mizoguchi, T. Takahashi, S. Aritome, and K. Takeuchi, "Data-Retention Characteristics Comparison of 2D and 3D TLC NAND Flash Memories," in *2017 IEEE International Memory Workshop (IMW)*, pp. 1-4.
- [6] Y. Luo, S. Ghose, Y. Cai, E. F. Haratsch, and O. Mutlu, "HeatWatch: Improving 3D NAND Flash Memory Device Reliability by Exploiting Self-Recovery and Temperature Awareness," in *2018 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pp. 504-517.
- [7] Q. Xiong, F. Wu, Z. Lu, Y. Zhu, Y. Zhou, Y. Chu, C. Xie, and P. Huang, "Characterizing 3D Floating Gate NAND Flash," in *2017 ACM SIGMETRICS / International Conference on Measurement and Modeling of Computer Systems*, pp. 31-32.
- [8] Y. Luo, S. Ghose, Y. Cai, E. F. Haratsch, and O. Mutlu, "Improving 3D NAND Flash Memory Lifetime by Tolerating Early Retention Loss and Process Variation," in *2018 ACM SIGMETRICS / International Conference on Measurement and Modeling of Computer Systems*, p. 106.
- [9] N. Mielke, H. P. Belgal, A. Fazio, Q. Meng, and N. Righos, "Recovery Effects in the Distributed Cycling of Flash Memories," in *2006 IEEE International Reliability Physics Symposium Proceedings*, pp. 29-35.
- [10] R. Micheloni, A. Marelli, and K. Eshghi, *Inside Solid State Drives (SSDs)*, ser. Springer Series in Advanced Microelectronics. Springer, 2013.
- [11] T. Parnell, N. Papandreou, T. Mittelholzer, and H. Pozidis, "Performance of Cell-to-cell Interference Mitigation in 1y-nm MLC Flash Memory," in *2015 Non-Volatile Memory Technology Symposium (NVMTS)*, pp. 1-4.
- [12] N. Papandreou, T. Parnell, T. Mittelholzer, H. Pozidis, T. Griffin, G. Tressler, T. Fisher, and C. Camp, "Effect of Read Disturb on Incomplete Blocks in MLC NAND Flash Arrays," in *2016 IEEE International Memory Workshop (IMW)*, pp. 1-4.