

Read Disturb Errors in MLC NAND Flash Memory

Yu Cai¹

Yixin Luo¹

Saugata Ghose¹

Erich F. Haratsch²

Ken Mai¹

Onur Mutlu^{3,1}

¹Carnegie Mellon University

²Seagate Technology

³ETH Zürich

This paper summarizes our work on experimentally characterizing, mitigating, and recovering read disturb errors in multi-level cell (MLC) NAND flash memory, which was published in DSN 2015 [16], and examines the work’s significance and future potential. NAND flash memory reliability continues to degrade as the memory is scaled down and more bits are programmed per cell. A key contributor to this reduced reliability is read disturb, where a read to one row of cells impacts the threshold voltages of unread flash cells in different rows of the same block. Such disturbances may shift the threshold voltages of these unread cells to different logical states than originally programmed, leading to read errors that hurt endurance.

For the first time in open literature, this work experimentally characterizes read disturb errors on state-of-the-art 2Y-nm (i.e., 20-24 nm) MLC NAND flash memory chips. Our findings (1) correlate the magnitude of threshold voltage shifts with read operation counts, (2) demonstrate how program/erase cycle count and retention age affect the read-disturb-induced error rate, and (3) identify that lowering pass-through voltage levels reduces the impact of read disturb and extend flash lifetime. Particularly, we find that the probability of read disturb errors increases with both higher wear-out and higher pass-through voltage levels.

We leverage these findings to develop two new techniques. The first technique mitigates read disturb errors by dynamically tuning the pass-through voltage on a per-block basis. Using real workload traces, our evaluations show that this technique increases flash memory endurance by an average of 21%. The second technique recovers from previously-uncorrectable flash errors by identifying and probabilistically correcting cells susceptible to read disturb errors. Our evaluations show that this recovery technique reduces the raw bit error rate by 36%.

1. Introduction

NAND flash memory currently sees widespread usage as a storage device, having been incorporated into systems ranging from mobile devices and client computers to data center storage, as a result of its increasing capacity and decreasing cost per bit. The increasing capacity and lower cost are mainly driven by aggressive transistor scaling and *multi-level cell* (MLC) technology, where a single flash cell can store more than one bit of data. However, as NAND flash memory capacity increases, flash memory suffers from different types of circuit-level noise, which greatly impact its reliability. These include program/erase cycling noise [8, 9], cell-to-cell program interference noise [8, 11, 14], retention noise [8, 10, 12, 13, 49, 59], and read disturb noise [19, 26, 59, 87].

Among all of these types of noise, *read disturb* noise has largely been understudied in the past for MLC NAND flash, with no open-literature work available prior to our DSN 2015 paper [16] that characterizes and analyzes the read disturb phenomenon.

One reason for this prior neglect has been the heretofore low occurrence of read-disturb-induced errors in older flash technologies. In *single-level cell* (SLC) NAND flash, read disturb errors were only expected to appear after an average of one million reads to a single flash block [26, 54]. Even with the introduction of MLC NAND flash, first-generation MLC devices were expected to exhibit read disturb errors after 100,000 reads [29, 54]. As a result of manufacturing process technology scaling, some modern MLC NAND flash devices are now prone to read disturb errors after as few as 20,000 reads, with this number expected to drop even further with continued scaling [29, 54]. The exposure of these read disturb errors can be exacerbated by the uneven distribution of reads across flash blocks in contemporary workloads [65, 89], where certain flash blocks experience high temporal locality and can, therefore, more rapidly exceed the read count at which read disturb errors are induced. We refer the reader to our prior works for a more detailed background [4, 5, 6, 16].

Read disturb errors are an intrinsic result of the flash architecture. Inside each flash cell, data is stored as the *threshold voltage* of the cell, based on the logical value that the cell represents. As shown in Figure 1, during a read operation to the cell, a *read reference voltage* (i.e., V_a , V_b , or V_c) is applied to the transistor corresponding to this cell. If this read reference voltage is higher than the threshold voltage of the cell, the transistor is turned *on*. The region in which the threshold voltage of a flash cell falls represents the cell’s current state, which can be ER (or erased), P1, P2, or P3. Each state decodes into a 2-bit value that is stored in the flash cell (e.g., 11, 10, 00, or 01). Note that the threshold voltage of all flash cells in a chip is bounded by an upper limit, V_{pass} , which is the *pass-through voltage*. More detailed explanations of how NAND flash memory cells work and the data retention errors in NAND flash memory can be found in our prior works [4, 5, 6, 10].

Within a flash block, the transistors of multiple cells, each from a different flash page, are tied together as a single *bitline*, which is connected to a single output wire. Only one cell is read at a time per bitline. In order to read one cell (i.e., to determine whether it is turned *on* or *off*), the transistors for the cells *not being read* must be kept *on* to allow the value from the cell being read to propagate to the output. This requires

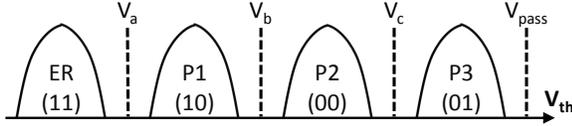


Figure 1: Threshold voltage distribution in 2-bit MLC NAND flash. Stored data values are represented as the tuple (LSB, MSB). Reproduced from [16].

the transistors to be powered with a *pass-through voltage*, which is a read reference voltage guaranteed to be higher than *any* stored threshold voltage (see Figure 1). Though these other cells are *not* being read, this high pass-through voltage induces electric tunneling that can shift the threshold voltages of these unread cells to higher values, thereby *disturbing the cell contents on a read operation to a neighboring page*. As we scale down the size of flash cells, the transistor oxide becomes thinner, which in turn increases this tunneling effect. With each read operation having an increased tunneling effect, it takes fewer read operations to neighboring pages for the unread flash cells to become disturbed (i.e., shifted to higher threshold voltages) and move into a different logical state.

In light of the increasing sensitivity of flash memory to read disturb errors, our goal is to (1) develop a thorough understanding of read disturb errors in state-of-the-art MLC NAND flash memories, by performing experimental characterization of such errors on existing commercial 2Y-nm (i.e., 20-24 nm) flash memory chips, and (2) develop mechanisms that can tolerate read disturb errors, making use of insights gained from our read disturb error characterization. The *key findings* from our quantitative characterization are:

- The effect of read disturb on threshold voltage distributions and raw bit error rates increases with both the number of reads to neighboring pages and the number of program/erase cycles on a block.
- Cells with lower threshold voltages are more susceptible to errors as a result of read disturb.
- As the pass-through voltage decreases, (1) the read disturb effect of each individual read operation becomes smaller, but (2) the read errors can increase due to reduced ability in allowing the read value to pass through the unread cells.
- If a page is recently written, a significant margin within the *ECC correction capability* (i.e., the total number of bit errors it can correct for a single read) is unused (i.e., the page can still tolerate more errors), which enables the page’s pass-through voltage to be lowered safely).

We exploit these studies on the relation between the read disturb effect and the pass-through voltage (V_{pass}), to design two mechanisms that reduce the reliability impact of read disturb. First, we propose a low-cost dynamic mechanism called *V_{pass} Tuning*, which, for each block, finds the lowest pass-through voltage that retains data correctness. *V_{pass} Tuning* extends flash endurance by exploiting the finding that a lower V_{pass} reduces the read disturb error count. Our evaluations using real workload traces show that *V_{pass} Tuning* extends

flash lifetime by 21%. Second, we propose *Read Disturb Recovery* (RDR), a mechanism that exploits the differences in the susceptibility of different cells to read disturb to extend the effective correction capability of error-correcting codes (ECC). RDR probabilistically identifies and corrects cells susceptible to read disturb errors. Our evaluations show that RDR reduces the raw bit error rate by 36%.

2. Characterizing Read Disturb in Real NAND Flash Memory Chips

We use an FPGA-based NAND flash testing platform in order to characterize read disturb on state-of-the-art flash chips [4,5,6,7]. We use the *read-retry* operation present within MLC NAND flash devices to accurately read the cell threshold voltage [4, 5, 6, 9, 10, 11, 12, 14, 15, 22, 69]. As threshold voltage values are proprietary information, we present our results using a *normalized threshold voltage*, where the nominal value of V_{pass} is equal to 512 in our normalized scale, and where 0 represents GND.

One limitation of using commercial flash devices is the inability to alter the V_{pass} value, as no such interface currently exists. We work around this by using the read-retry mechanism, which allows us to change the read reference voltage V_{ref} one wordline at a time. Since both V_{pass} and V_{ref} are applied to wordlines, we can mimic the effects of changing V_{pass} by instead changing V_{ref} and examining the impact on the wordline being read. We perform these experiments on one wordline per block, and repeat them over ten different blocks.

We present our major findings below. For a complete description of all of our observations, we refer the reader to our DSN 2015 paper [16].

2.1. Quantifying Read Disturb Perturbations

First, we quantify the amount by which read disturb shifts the threshold voltage, by measuring threshold voltage values for unread cells after 0, 250K, 500K, and 1 million read operations to other cells within the same flash block. Figure 2a shows the distribution of the threshold voltages for cells in a flash block after 0, 250K, 500K, and 1 million read operations. Figure 2b zooms in on this to illustrate the distribution for values in the ER state. We find that *the magnitude of the threshold voltage shift for a cell due to read disturb (1) increases with the number of read disturb operations, and (2) is higher if the cell has a lower threshold voltage*.

2.2. Effect of Read Disturb on Raw Bit Error Rate

Second, we aim to relate these threshold voltage shifts to the *raw bit error rate* (RBER), which refers to the probability of reading an incorrect state from a flash cell. We measure whether flash cells that are more worn out (i.e., cells that have been programmed and erased more times) are impacted differently due to read disturb. Figure 3 shows the RBER over an increasing number of read disturb operations for different

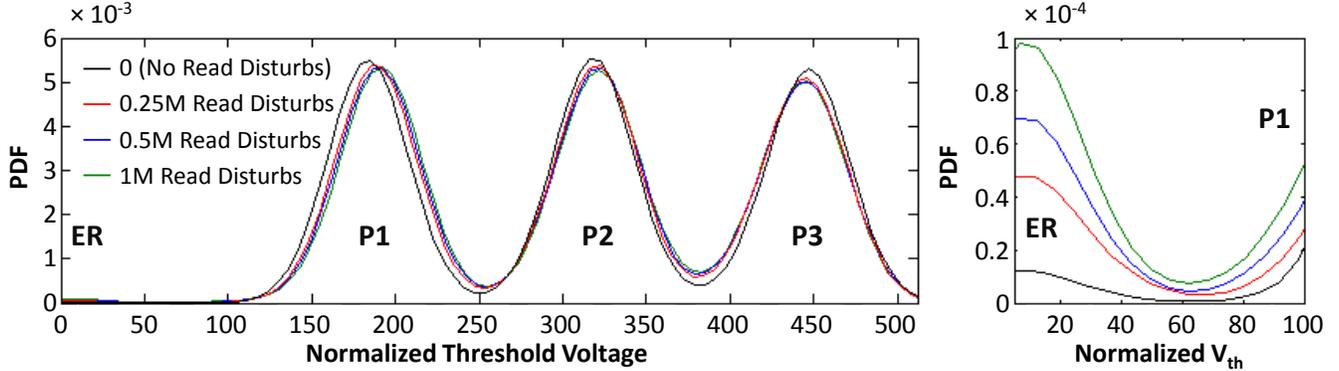


Figure 2: (a) Threshold voltage distribution of all programmed states before and after read disturb; (b) Threshold voltage distribution between erased state and P1 state. Reproduced from [16].

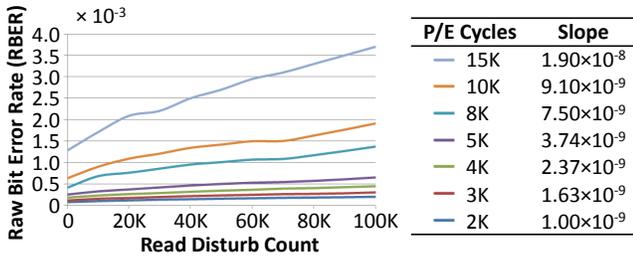


Figure 3: Raw bit error rate vs. read disturb count under different levels of program and erase (P/E) wear. Reproduced from [16].

amounts of P/E cycle wear (i.e., the amount of wearout in P/E cycles) on flash blocks. Each level shows a linear RBER increase as the read disturb count increases. We find that (1) for a given amount of P/E cycle wear on a block, the raw bit error rate increases roughly linearly with the number of read disturb operations, and that (2) the effects of read disturb are greater for cells that have experienced a larger number of P/E cycles.

2.3. Pass-Through Voltage Impact on Read Disturb

Third, we show that the cause of read disturb can be reduced by reducing (i.e., relaxing) the pass-through voltage using a circuit-level model of the flash cell, and verify this observation using real measurements. Figure 4 shows the measured change in RBER as a function of the number of read operations, for selected relaxations of V_{pass} . Note that the x-axis uses a log scale. For a fixed number of reads, even a small decrease in the V_{pass} value can yield a significant decrease in RBER. As an example, at 100K reads, lowering V_{pass} by 2% can reduce the RBER by as much as 50%. Conversely, for a fixed RBER, a decrease in V_{pass} exponentially increases the number of tolerable read disturbs. However, decreasing V_{pass} can prevent some cells' values from propagating correctly along the bitline on a read, as an unread flash cell transistor may be incorrectly turned off, thus generating new errors.

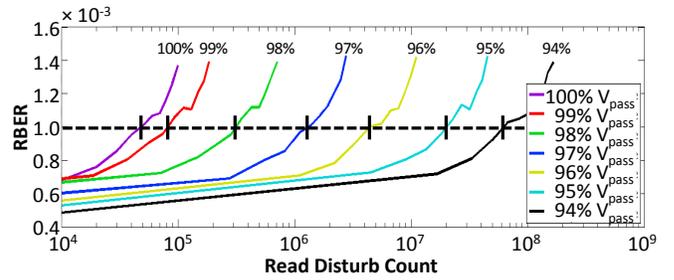


Figure 4: Raw bit error rate vs. read disturb count for different V_{pass} values, for flash memory under 8K program/erase cycles of wear. Reproduced from [16].

Unlike read disturb errors, these bitline propagation errors (or read errors) do not alter the threshold voltage of the flash cell.

2.4. Effect of Pass-Through Voltage on Raw Bit Error Rate

Fourth, setting V_{pass} to a value slightly lower than the maximum V_{th} leads to a trade-off. On the one hand, it can substantially reduce the effects of read disturb. On the other hand, it causes a small number of unread cells to incorrectly stay off instead of passing through a value, potentially leading to a read error. Therefore, if the number of read disturb errors can be dropped significantly by lowering V_{pass} , the small number of read errors introduced may be warranted. If too many read errors occur, we can always fall back to using the maximum threshold voltage for V_{pass} without consequence. Naturally, this trade-off depends on the magnitude of these error rate changes. We now explore the gains and costs, in terms of overall RBER, for relaxing V_{pass} below the maximum threshold voltage of a block.

To identify the extent to which relaxing V_{pass} affects the raw bit error rate, we experimentally sweep over V_{pass} , re-reading the data after a range of different retention ages, as shown in Figure 5. First, we observe that across all of our studied retention ages, V_{pass} can be lowered to some degree without inducing any read errors. For greater relaxations, though,

the error rate increases as more unread cells are incorrectly turned *off* during read operations. We also note that, for a given V_{pass} value, the additional read error rate is lower if the read is performed a longer time after the data is programmed into the flash (i.e., if the retention age is longer). This is because of the retention loss effect, where cells slowly leak charge and thus have lower threshold voltage values over time. Naturally, as the threshold voltage of every cell decreases, a relaxed V_{pass} becomes more likely to correctly turn *on* the unread cells.

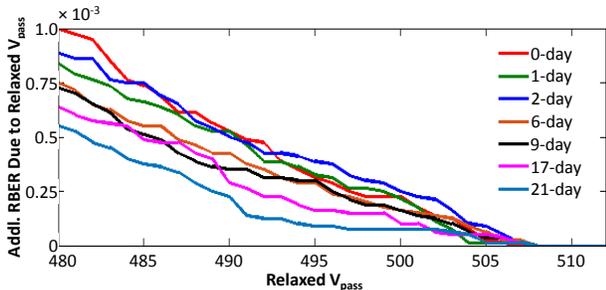


Figure 5: Additional raw bit error rate induced by relaxing V_{pass} , shown across a range of data retention ages. Reproduced from [16].

2.5. Error Correction with Reduced Pass-Through Voltage

Fifth, while we have shown, in Section 3.6 of our DSN 2015 paper [16], that V_{pass} can be lowered to some degree without introducing new raw bit errors, we would ideally like to further decrease V_{pass} to lower the read disturb impact more. This can enable flash devices to tolerate many more reads. The ECC used for NAND flash memory can tolerate an RBBER of up to 10^{-3} [12, 13], which occurs only during worst-case conditions such as long retention time. Our goal is to identify how many additional raw bit errors the current level of ECC provisioning in flash chips can sustain. Figure 6 shows how the expected RBBER changes over a 21-day period for our tested flash chip *without read disturb*, using a block with 8,000 P/E cycles of wear. An RBBER margin (20% of the total ECC correction capability) is reserved to account for variations in the distribution of errors and other potential errors (e.g., program and erase errors). For each retention age, the maximum percentage of *safe* V_{pass} reduction (i.e., the lowest value of V_{pass} at which all read errors can still be corrected by ECC) is listed on the top of Figure 6. As we can see, by exploiting the previously-unused ECC correction capability, V_{pass} can be safely reduced by as much as 4% when the retention age is low (less than 4 days).

Our key insight from this study is that a lowered V_{pass} can reduce the effects of read disturb, and that the read errors induced from lowering V_{pass} can be tolerated by the built-in error correction mechanism within modern flash controllers. More results and more detailed analysis are in our DSN 2015 paper [16].

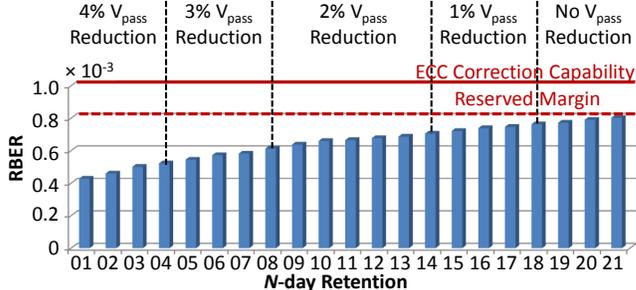


Figure 6: Overall raw bit error rate and tolerable V_{pass} reduction vs. retention age, for a flash block with 8K P/E cycles of wear. Reproduced from [16].

3. Mitigation: Pass-Through Voltage Tuning

To minimize the effect of read disturb, we propose a mechanism called V_{pass} Tuning, which *learns the minimum pass-through voltage for each block, such that all data within the block can be read correctly with ECC*. Figure 7 provides an exaggerated illustration of how the unused ECC capability changes over the retention period (i.e., the *refresh interval*). At the start of each retention period, there are no retention errors or read disturb errors, as the data has just been restored. In these cases, the large unused ECC capability allows us to design an *aggressive* read disturb mitigation mechanism, as we can safely *introduce correctable errors*. Thanks to read disturb mitigation, we can reduce the effect of each individual read disturb, thus lowering the total number of read disturb errors accumulated by the end of the refresh interval. This reduction in read disturb error count leads to lower *error count peaks* at the end of each refresh interval, as shown in Figure 7 by the distance between the solid black line and the dashed red line. Since flash lifetime is dictated by the number of data errors (i.e., when the total number of errors exceeds the ECC correction capability, the flash device has reached the end of its life), lowering the error count peaks extends lifetime by extending the time before these peaks exhaust the ECC correction capability.

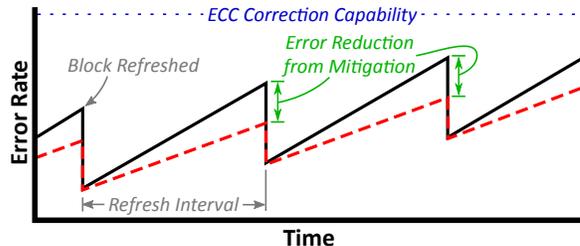


Figure 7: Exaggerated example of how read disturb mitigation reduces error rate peaks for each refresh interval. Solid black line is the unmitigated error rate, and dashed red line is the error rate after mitigation. (Note that the error rate does not include read errors introduced by reducing V_{pass} , as the unused error correction capability can tolerate errors caused by V_{pass} Tuning.) Reproduced from [16].

Our learning mechanism works online and is triggered on a daily basis. V_{pass} Tuning can be fully implemented within the *flash controller*, and has two components:

1. It first finds the size of the ECC margin M (i.e., the unused correction capability within ECC) that can be exploited to tolerate additional read errors for each block. In order to do this, our mechanism discovers the page with *approximately* the highest number of raw bit errors.
2. Once it knows the available margin M , our mechanism calibrates the pass-through voltage V_{pass} on a *per-block basis* to find the lowest value of V_{pass} that introduces no more than M additional raw errors.

The first component of our mechanism must first approximately discover the page with the highest error count, which we call the *predicted worst-case page*. After manufacturing, we statically find the predicted worst-case page by programming pseudo-randomly generated data to each page within the block, and then immediately reading the page to find the error count, as prior work on error analysis has done [8]. For each block, we record the page number of the page with the highest error count. Our mechanism obtains the error count, which we define as our *maximum estimated error (MEE)*, by performing *a single read* to this page and reading the error count provided by ECC (once a day). We conservatively reserve 20% of the spare ECC correction capability in our calculations. Thus, if the maximum number of raw bit errors correctable by ECC is C , we calculate the available ECC margin for a block as $M = (1 - 0.2) \times C - MEE$.

The second component of our mechanism identifies the greatest V_{pass} reduction that introduces no more than M raw bit errors. The general V_{pass} *identification process* requires three steps:

Step 1: Aggressively reduce V_{pass} to $V_{pass} - \Delta$, where Δ is the smallest resolution by which V_{pass} can change.

Step 2: Apply the new V_{pass} to *all* wordlines in the block. Count the number of 0's read from the page (i.e., the number of bitlines incorrectly switched *off*) as N . If $N \leq M$ (recall that M is the extra available ECC correction margin), the read errors resulting from this V_{pass} value can be corrected by ECC, so we repeat Steps 1 and 2 to try to further reduce V_{pass} . If $N > M$, it means we have reduced V_{pass} too aggressively, so we proceed to Step 3 to roll back to an acceptable value of V_{pass} .

Step 3: Increase V_{pass} to $V_{pass} + \Delta$, and verify that the introduced read errors can be corrected by ECC (i.e., $N \leq M$). If this verification fails, we repeat Step 3 until the read errors are reduced to an acceptable range.

The implementation can be simplified greatly in practice, as the error rate changes are relatively slow over time. Over the course of the seven-day refresh interval, our mechanism must perform one of two actions each day:

Action 1: When a block is *not* refreshed, our mechanism checks once daily if V_{pass} should *increase*, to accommodate the slowly-increasing number of errors due to dynamic factors (e.g., retention errors, read disturb errors).

Action 2: When a block is refreshed, all retention and read disturb errors accumulated during the previous refresh interval are corrected. At this time, our mechanism checks how much V_{pass} can be *lowered* by.

Our mechanism repeats the V_{pass} identification process for each block that contains valid data to learn the *minimum pass-through voltage we can use*. It also repeats the entire V_{pass} learning process daily to adapt to threshold voltage changes due to retention loss [11, 14]. As such, the pass-through voltage of *all blocks* in a flash drive can be fine-tuned *continuously* to reduce read disturb and thus improve overall flash lifetime. Our DSN 2015 paper [16] describes this mechanism in more detail, and discusses a fallback mechanism for extreme cases where the additional errors accumulating between tunings exceed our 20% margin of unused error correction capability. For more detail, we refer the reader to Section 4 of our DSN 2015 paper [16].

Our mechanism can reduce V_{pass} by as much as 4%. Through a series of optimizations, described in more detail in Section 4 of our DSN 2015 paper [16], it only incurs an average daily performance overhead of 24.34 sec for a 512GB SSD, and uses only 128KB storage overhead to record per-block data.

We evaluate V_{pass} *Tuning* with I/O traces collected from a wide range of real workloads with different use cases [38, 43, 65, 83, 89]. Figure 8 shows how our mechanism can increase the endurance (measured as the number of program/erase cycles that take place before the NAND flash memory can no longer be used). We find that for a variety of our workloads, our V_{pass} tuning mechanism increases flash memory endurance by an average of 21.0%, thanks to its success in reducing the number of raw bit errors that occur due to read disturb.

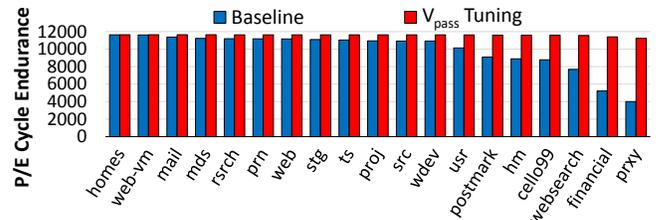


Figure 8: Endurance improvement with V_{pass} *Tuning*. Reproduced from [16].

4. Read Disturb Oriented Error Recovery

Even if we mitigate the impact of read disturb errors, a flash device will eventually exhaust its lifetime. At that point, some reads will have more raw errors to correct than can be corrected by ECC, preventing the drive from returning the correct data to the user. Traditionally, this is referred to as the point of *data loss*.

We propose to take advantage of our understanding of read disturb behavior, by designing a mechanism that can recover data *even after the device has exceeded its lifetime*.

This mechanism, which we call *Read Disturb Recovery* (RDR), (1) identifies flash cells that are *susceptible* to generating errors due to read disturb (i.e., *disturb-prone* cells), and (2) probabilistically corrects the data stored in these cells without the assistance of ECC. After these probabilistic corrections, the number of errors for a read will be brought back down, to a point at which ECC can successfully correct the remaining errors and return valid data to the user.

To understand why identifying disturb-prone cells can help with correcting errors, we study why read disturb errors occur to begin with. Figure 9a shows the state of four flash cells before read disturb happens. The two blue cells are both programmed with a two-bit value of 11, and the two red cells are programmed with a two-bit value of 00, with each two-bit value being assigned to a different range of threshold voltages (V_{th}). Between each assigned range is a margin. When read disturb occurs, the blue cells, which are *disturb-prone*, experience large V_{th} shifts upwards, while the blue cells, which are *disturb-resistant*, do not shift much, as shown in Figure 9b. Now that the distributions of these two-bit values overlap, a read error will occur for these four cells.

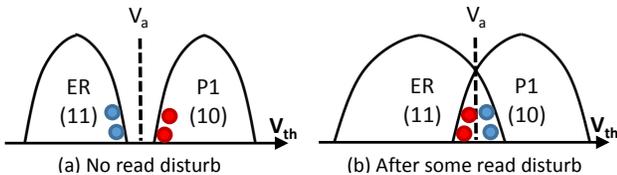


Figure 9: V_{th} distributions before and after read disturb. Reproduced from [16].

Identifying Susceptible Cells. In order to identify susceptible cells, RDR induces a significant number of additional read disturbs (e.g., 100K) within the flash cells that contain uncorrectable errors. We do this by characterizing the degree of the threshold voltage shift (ΔV_{th}) induced by the additional read disturbs, and comparing the shift to a delta threshold voltage (ΔV_{ref}) at the intersection of the two probability density functions. We classify cells with a higher threshold voltage change ($\Delta V_{th} > \Delta V_{ref}$) as *disturb-prone* cells. We classify cells with a lower or negative threshold voltage change ($\Delta V_{th} < \Delta V_{ref}$) as *disturb-resistant* cells. Section 5.2 of our DSN 2015 paper [16] provides more detailed results and analysis of disturb-prone and disturb-resistant cells.

Correcting Susceptible Cells. For flash cells with threshold voltages close to the boundary between two different data values, RDR predicts that the disturb-prone cells belong to the lower of the two voltage distributions (ER in Figure 9). Likewise, disturb-resistant cells near the boundary likely belong to the higher voltage distribution (P1 in Figure 9). This does *not* eliminate all errors, but decreases the raw bit errors in disturb-prone cells. RDR attempts to correct the remaining raw bit errors using ECC. Section 5.3 of our DSN 2015 paper [16] provides more detail on the RDR mechanism.

We evaluate how the overall RBER changes when we use RDR. Fig. 10 shows experimental results for error recovery in a flash block with 8,000 P/E cycles of wear. When RDR is applied, the *reduction in overall RBER grows with the read disturb count, from a few percent for low read disturb counts up to 36% for 1 million read disturb operations*. As data experiences a greater number of read disturb operations, the read disturb error count contributes to a significantly larger portion of the total error count, which our recovery mechanism targets and reduces. We therefore conclude that RDR can provide a large effective extension of the ECC correction capability.

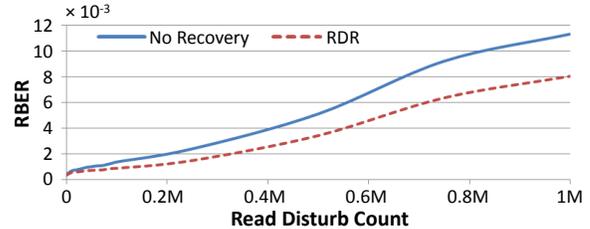


Figure 10: Raw bit error rate vs. number of read disturb operations, with and without RDR, for a flash block with 8,000 P/E cycles of wear. Reproduced from [16].

5. Related Work

We break down related work on NAND flash memory (Section 5.1) into six major categories: (1) read disturb error characterization, (2) NAND flash memory error characterization, (3) 3D NAND error characterization, (4) read disturb error mitigation, (5) voltage optimization, and (6) error recovery. We then introduce related work on read disturb errors in DRAM (Section 5.2) and emerging memory technologies (Section 5.3).

5.1. Related Works on NAND Flash Memory

Read Disturb Error Characterization. Prior to this work [16], the read disturb phenomenon for NAND flash memory has not been well explored in openly-available literature. Prior work [42] experimentally characterizes and proposes solutions for read disturb errors in DRAM. The mechanisms for disturbance and techniques to mitigate them are different between DRAM and NAND flash due to device-level differences [61]. Recent work has characterized concentrated read disturb effect and find that there are more read disturb errors on the direct neighbors to the page being repeatedly read [97]. Recent work has found that read disturb errors significantly reduce the reliability of unprogrammed and partially-programmed wordlines within a flash block, and can cause security vulnerabilities [15, 67]. These unprogrammed and partially programmed wordlines have lower threshold voltages (e.g., all cells in unprogrammed wordlines are in erased state), they are more sensitive to read disturb effect. When the wordlines are fully-programmed, NAND flash memory chip cannot correct any of these read disturb errors and thus program the misread flash cells into an incorrect state.

NAND Flash Memory Error Characterization. There are many past works from us and other research groups that analyze many different types of NAND flash memory errors in MLC, planar NAND flash memory, including P/E cycling errors [9, 52, 59, 68, 72], programming errors [15, 52, 72], cell-to-cell program interference errors [9, 11, 14], retention errors [9, 10, 12, 59, 68], and read disturb errors [16, 59, 68], and propose many different mitigation mechanisms. These works complement our DSN 2015 paper. A survey of these works (and many other related ones) can be found in our recent works [4, 5, 6]. These works characterize how raw bit error rate and threshold voltage change over various types of noise. Our recent work characterizes the same types of errors in TLC, planar NAND flash memory and has similar findings [4, 5, 6]. Thus, we believe that most of the findings on MLC NAND flash memory can be generalized to any types of planar NAND flash memory devices (e.g., SLC, MLC, TLC, or QLC). Recent work has also studied SSD errors in the field, and has shown the system-level implications of these errors to large-scale data centers [56, 66, 77].

3D NAND Error Characterization. Recently, manufacturers have begun to produce SSDs that contain *three-dimensional* (3D) NAND flash memory [33, 37, 57, 58, 70, 96]. In 3D NAND flash memory, *multiple layers* of flash cells are stacked vertically to increase the density and to improve the scalability of the memory [96]. In order to achieve this stacking, manufacturers have changed a number of underlying properties of the flash memory design. However, the internal organization of a flash block remains unchanged. Thus, read disturb errors are similar in 3D NAND flash memory. But the rate of read disturb errors are significantly reduced in today’s 3D NAND because it currently uses a larger manufacturing process technology [23, 25]. We refer the reader to our prior work for a more detailed comparison between 3D NAND and planar NAND [4, 5, 6]. Recent work characterizes the latency and raw bit error rate of 3D NAND devices based on floating gate cells [94] and make similar observations as in planar NAND devices based on floating gate cells. Recent work has reported several differences between 3D NAND and planar NAND through circuit level measurements. These differences include 1) smaller program variation at high P/E cycle [70], 2) smaller program interference [70], 3) early retention loss [17, 17, 60]. We characterize the impact of dwell time, i.e., idle time between consecutive program cycles, and environment temperature on the retention loss speed and programming accuracy in 3D charge trap NAND flash cells [53]. The field (both academia and industry) is currently in much need of rigorous experimental characterization and analysis of 3D NAND flash memory devices.

Read Disturb Error Mitigation. Prior work proposes to mitigate read disturb errors by caching recently read data to avoid a read operation [85]. Prior work also proposes to mitigate read disturb errors using an idea similar to remapping-based refresh [12], known as *read reclaim*. The key idea

of read reclaim is to remap the data in a block to a new flash block, if the block has experienced a high number of reads [21, 29, 30, 40]. To bound the number of read disturb errors, some flash vendors specify a maximum number of tolerable reads for a flash block, at which point read reclaim rewrites the data to a new block (just as is done for remapping-based refresh).

Two mechanisms are currently being implemented within Yaffs (Yet Another Flash File System) to handle read disturb errors, though they are not yet available [54]. The first mechanism is similar to read reclaim [29], where a block is rewritten after a fixed number of page reads are performed to the block (e.g., 50,000 reads for an MLC chip). The second mechanism periodically inserts an additional read (e.g., a read every 256 block reads) to a page within the block, to check whether that page has experienced a read disturb error, in which case the page is copied to a new block.

Recent work proposes to remap read-hot pages to blocks configured as SLC, which are resistant to read disturb [48, 100]. Ha et al. combine this read-hot page mapping technique with our V_{pass} Tuning technique and read reclaim [30] to further reduce read disturb errors. This shows that the techniques proposed by prior work are orthogonal to our read disturb mitigation techniques, and can be combined with our work for even greater protection.

Voltage Optimization. While the pass-through voltage optimization is specific to read disturb error mitigation, a few works that propose optimizing the read reference voltage have the same spirit [11, 14, 68]. Cai et al. propose a technique to calculate the optimal read reference voltage from the mean and variance of the threshold voltage distributions [14], which are characterized by the read-retry technique [9]. The cost of such a technique is relatively high, as it requires periodically reading flash memory with all possible read reference voltages to discover the threshold voltage distributions. Papandreou et al. propose to apply a per-block close-to-optimal read reference voltage by periodically sampling and averaging 6 OPTs within each block, learned by exhaustively trying all possible read reference voltages [68]. In contrast, ROR can find the actual optimal read reference voltage at a much lower latency, thanks to the new findings and observations in our DSN 2015 paper [10]. We already showed in our DSN 2015 paper that ROR greatly outperforms naive read-retry, which is significantly simpler than the mechanism proposed in [68].

Recently, Luo et al. propose to accurately predict the optimal read reference voltage using an online flash channel model for each chip learned online [52]. Cai et al. proposes a new technique called V_{pass} tuning, which tunes the *pass-through voltage*, i.e., a high reference voltage applied to turn on unread cells in a block, to mitigate read disturb errors [16]. Du et al. proposes to tune the optimal read reference voltages for ECC code soft decoding to improve ECC correction capability [20]. Fukami et al. proposes to use read-retry to

improve the reliability of chip-off forensic analysis of NAND flash memory devices [22].

Error Recovery. To our knowledge, no prior work other than our DSN 2015 paper can recover the data from an uncorrectable error that is beyond the error correction capability of ECC caused by read disturb [16]. We have proposed a mechanism called RFR to opportunistically recover from uncorrectable data retention errors [4, 5, 6, 16]. RFR, similar to RDR proposed in this work, identifies *fast- and slow-leaking cells*, rather than disturb-prone and disturb-resistant cells, and probabilistically correct uncorrectable retention errors offline.

5.2. Read Disturb Errors in DRAM

Commodity DRAM chips that are sold and used in the field today exhibit read disturb errors [42], also called *RowHammer*-induced errors [61], which are *conceptually* similar to the read disturb errors found in NAND flash memory. Repeatedly accessing the same row in DRAM can cause bit flips in data stored in adjacent DRAM rows. In order to access data within DRAM, the row of cells corresponding to the requested address must be *activated* (i.e., opened for read and write operations). This row must be *precharged* (i.e., closed) when another row in the same DRAM bank needs to be activated. Through experimental studies on a large number of real DRAM chips, we show that when a DRAM row is activated and precharged repeatedly (i.e., *hammered*) enough times within a DRAM refresh interval, one or more bits in physically-adjacent DRAM rows can be flipped to the wrong value [42].

We tested 129 DRAM modules manufactured by three major manufacturers (A, B, and C) between 2008 and 2014, using an FPGA-based experimental DRAM testing infrastructure [31] (more detail on our experimental setup, along with a list of all modules and their characteristics, can be found in our original RowHammer paper [42]). Figure 11 shows the rate of RowHammer errors that we found, with the 129 modules that we tested categorized based on their manufacturing date. We find that 110 of our tested modules exhibit RowHammer errors, with the earliest such module dating back to 2010. In particular, we find that *all* of the modules manufactured in 2012–2013 that we tested are vulnerable to RowHammer. Like with many NAND flash memory error mechanisms, especially read disturb, RowHammer is a recent phenomenon that especially affects DRAM chips manufactured with more advanced manufacturing process technology generations.

Figure 12 shows the distribution of the number of rows (plotted in log scale on the y-axis) within a DRAM module that flip the number of bits along the x-axis, as measured for example DRAM modules from three different DRAM manufacturers [42]. We make two observations from the figure. First, the number of bits flipped when we hammer a row (known as the *aggressor row*) can vary significantly within a module. Second, each module has a different distribution

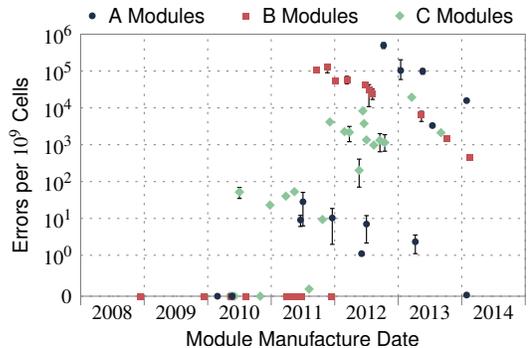


Figure 11: RowHammer error rate vs. manufacturing dates of 129 DRAM modules we tested. Reproduced from [42].

of the number of rows. Despite these differences, we find that this DRAM failure mode affects more than 80% of the DRAM chips we tested [42]. As indicated above, this read disturb error mechanism in DRAM is popularly called RowHammer [61].

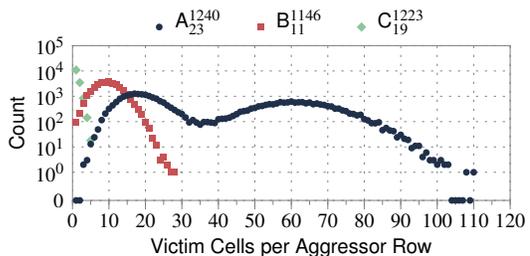


Figure 12: Number of victim cells (i.e., number of bit errors) when an aggressor row is repeatedly activated, for three representative DRAM modules from three major manufacturers. We label the modules in the format X_n^{yyww} , where X is the manufacturer (A, B, or C), $yyww$ is the manufacture year (yy) and week of the year (ww), and n is the number of the selected module. Reproduced from [42].

Various recent works show that RowHammer can be maliciously exploited by user-level software programs to (1) induce errors in existing DRAM modules [42, 61] and (2) launch attacks to compromise the security of various systems [2, 3, 27, 28, 34, 61, 74, 76, 78, 79, 90, 93]. For example, by exploiting the RowHammer read disturb mechanism, a user-level program can gain kernel-level privileges on real laptop systems [78, 79], take over a server vulnerable to RowHammer [28], take over a victim virtual machine running on the same system [2], and take over a mobile device [90]. Thus, the RowHammer read disturb mechanism is a prime (and perhaps the first) example of how a circuit-level failure mechanism in DRAM can cause a practical and widespread system security vulnerability.

Note that various solutions to RowHammer exist [41, 42, 61], but we do not discuss them in detail here. Our recent work [61] provides a comprehensive overview. A very promising proposal is to modify either the memory controller or the DRAM chip such that it probabilistically refreshes the physically-adjacent rows of a recently-activated row, with very low probability. This solution is called *Probabilistic Ad-*

adjacent Row Activation (PARA) [42]. Our prior work shows that this low-cost, low-complexity solution, which does not require any storage overhead, greatly closes the RowHammer vulnerability [42].

The RowHammer effect in DRAM worsens as the manufacturing process scales down to smaller node sizes [42,61,62,63]. More findings on RowHammer, along with extensive experimental data from real DRAM devices, can be found in our prior works [41, 42, 61].

5.3. Errors in Emerging Memory Technologies

Emerging nonvolatile memories [55], such as *phase-change memory* (PCM) [45, 46, 47, 75, 92, 95, 99], *spin-transfer torque magnetic RAM* (STT-RAM or STT-MRAM) [44, 64], *metal-oxide resistive RAM* (RRAM) [91], and *memristors* [18, 84], are expected to bridge the gap between DRAM and NAND-flash-memory-based SSDs, providing DRAM-like access latency and energy, and at the same time SSD-like large capacity and nonvolatility (and hence SSD-like data persistence). While their underlying designs are different from DRAM and NAND flash memory, these emerging memory technologies have been shown to exhibit similar types of errors.

PCM-based devices are expected to have a limited lifetime, as PCM can only endure a certain number of writes [45, 75, 92], similar to the P/E cycling errors in SSDs (though PCM's write endurance is higher than that of SSDs). PCM suffers from (1) *resistance drift* [32, 73, 92], where the resistance used to represent the value becomes higher over time (and eventually can introduce a bit error), similar to how charge leakage in NAND flash memory and DRAM lead to retention errors over time; and (2) *write disturb* [35], where the heat generated during the programming of one PCM cell dissipates into neighboring cells and can change the value that is stored within the neighboring cells, similar in concept to cell-to-cell program interference in NAND flash memory.

STT-RAM suffers from (1) *retention failures*, where the value stored for a single bit (as the magnetic orientation of the layer that stores the bit) can flip over time [36, 82, 86]; and (2) *read disturb* (a conceptually different phenomenon from the read disturb in DRAM and flash memory), where reading a bit in STT-RAM can inadvertently induce a write to that same bit [64].

Due to the nascent nature of emerging nonvolatile memory technologies and the lack of availability of large-capacity devices built with them, extensive and dependable experimental studies have yet to be conducted on the reliability of real PCM, STT-RAM, RRAM, and memristor chips. However, we believe that error mechanisms conceptually or abstractly similar to those for flash memory and DRAM are likely to be prevalent in emerging technologies as well (as supported by some recent studies [1, 35, 39, 64, 80, 81, 98]), albeit with different underlying mechanisms and error rates.

6. Significance

Our DSN 2015 paper [16] is the first openly-available work to (1) characterize the impact of read disturb errors on commercially-available NAND flash memory devices, and (2) propose novel solutions to the read disturb errors that minimize them or recover them after error occurrence. We believe that our characterization results, analyses, and mechanisms can have a wide impact on future research on read disturb and NAND flash memory reliability.

6.1. Long-Term Impact

As flash devices continue to become more pervasive, there is renewed concern about the fewer number of writes that these flash devices can endure as they continue to scale [19, 29, 54]. This lower write endurance is a result of the larger number of errors introduced from manufacturing process technology scaling, and the use of multi-level cell technology. Today's planar NAND flash devices can endure only on the order of 100 program and erase cycles [71] without the assistance of aggressive error mitigation techniques such as data refresh [12, 50].

While there are several solutions for other types of NAND flash memory errors, read disturb has in the past been largely neglected because it has only become a significant problem at these smaller process technology nodes [29, 54]. Our work has the potential to change this relative lack of attention to read disturb for several reasons:

- We demonstrate on existing devices that read disturb is a significant problem today, and that it contributes a large number of errors that further reduce NAND flash memory endurance.
- We provide key insights as to *why* these errors occur, as well as why they will only worsen as technology scaling progresses.
- We show that it is possible to develop lightweight solutions that can alleviate the impact of read disturb.

Unfortunately, unless error mitigation techniques for read disturb are deployed in production NAND flash memory, read disturb will continue to negatively impact flash lifetime. While today's 3D NAND flash devices use larger process technologies that are less prone to read disturb effects [6, 51], future 3D NAND flash chips are expected to return to using smaller process technologies that remain susceptible to read disturb, as manufacturers continue to aggressively increase flash device densities [24, 88, 96]. With flash devices expected to remain a large component of the storage market for the foreseeable future, and with continued demand for higher flash densities, we expect that our work on read disturb can inspire manufacturers and researchers to adopt effective solutions to the read disturb problem.

The recovery mechanism that we propose, RDR, provides a new protective scheme for data storage that people have not considered before. Today, an increasingly larger volume of

data is stored in data centers belonging to cloud service providers, who must provide a strong guarantee of data integrity for their end users. With flash storage continuing to expand in data centers [56, 66, 77], RDR (as well as other recovery solutions that RDR might inspire) can reduce the probability of unrecoverable data loss for high-density storage. In fact, the availability of a recovery mechanism like RDR can also influence more data centers to adopt flash memory for storage.

6.2. New Research Directions

In our DSN 2015 paper [16], we present a number of new quantitative results on the impact of read disturb errors on NAND flash reliability, as well as how several key factors affect the number of errors induced by read disturb, such as the pass-through voltage, the number of program/erase cycles, and the retention age. Such a detailed characterization was not openly available in the past. We believe that by releasing our characterization data, researchers in both academia and industry will be able to use the data to develop further mechanisms for read disturb recovery and mitigation. In addition, by exposing the importance of the read disturb problem in contemporary NAND flash devices, we expect that our work will draw more attention to the problem, and will inspire other researchers to further characterize and understand the read disturb phenomenon.

In fact, one of our recent works builds on our DSN 2015 paper and shows that read disturb errors can potentially cause security vulnerabilities in modern SSDs [15].

We also expect that RDR, our recovery approach, will inspire researchers to design other data recovery mechanisms for NAND flash memory that also leverage the intrinsic properties of flash devices. To our knowledge, our new data recovery mechanism is the first to do so, by discovering and exploiting the variation in read disturb shifts that arise from the underlying process variation within a flash chip.

7. Conclusion

We provide the first detailed experimental characterization of read disturb errors for 2Y-nm MLC NAND flash memory chips. We find that bit errors due to read disturb are much more likely to take place in cells with lower threshold voltages, as well as in cells with greater wear. We also find that reducing the pass-through voltage can effectively mitigate read disturb errors. Using these insights, we propose (1) a mitigation mechanism, called V_{pass} Tuning, which dynamically adjusts the pass-through voltage for each flash block online to minimize read disturb errors, and (2) an error recovery mechanism, called *Read Disturb Recovery*, which exploits the differences in susceptibility of different cells to read disturb, to probabilistically correct read disturb errors. We hope that our characterization and analysis of the read disturb phenomenon enables the development of other error mitigation and tolerance mechanisms, which will become increasingly

necessary as continued flash memory scaling leads to greater susceptibility to read disturb. We also hope that our results will motivate NAND flash manufacturers to add pass-through voltage controls to next-generation chips, allowing flash controller designers to exploit our findings and design controllers that tolerate read disturb more effectively.

Acknowledgments

We thank the anonymous reviewers for their feedback. This work is partially supported by the Intel Science and Technology Center, the CMU Data Storage Systems Center, and NSF grants 0953246, 1065112, 1212962, and 1320531.

References

- [1] A. Athmanathan, M. Stanisavljevic, N. Papandreou, H. Pozidis, and E. Eleftheriou, "Multilevel-Cell Phase-Change Memory: A Viable Technology," *JETCAS*, 2016.
- [2] E. Bosman, K. Razavi, H. Bos, and C. Guiffrida, "Dedup Est Machina: Memory Deduplication as an Advanced Exploitation Vector," in *SP*, 2016.
- [3] W. Burleson, O. Mutlu, and M. Tiwari, "Who is the Major Threat to Tomorrow's Security? You, the Hardware Designer," in *DAC*, 2016.
- [4] Y. Cai, S. Ghose, E. F. Haratsch, Y. Luo, and O. Mutlu, "Error Characterization, Mitigation, and Recovery in Flash-Memory-Based Solid-State Drives," *Proc. IEEE*, Sep. 2017.
- [5] Y. Cai, S. Ghose, E. F. Haratsch, Y. Luo, and O. Mutlu, "Error Characterization, Mitigation, and Recovery in Flash Memory Based Solid-State Drives," arXiv:1706.08642 [cs.AR], 2017.
- [6] Y. Cai, S. Ghose, E. F. Haratsch, Y. Luo, and O. Mutlu, "Errors in Flash-Memory-Based Solid-State Drives: Analysis, Mitigation, and Recovery," arXiv:1711.11427 [cs.AR], 2017.
- [7] Y. Cai, E. F. Haratsch, M. P. McCartney, and K. Mai, "FPGA-Based Solid-State Drive Prototyping Platform," in *FCCM*, 2011.
- [8] Y. Cai, E. F. Haratsch, O. Mutlu, and K. Mai, "Error Patterns in MLC NAND Flash Memory: Measurement, Characterization, and Analysis," in *DATE*, 2012.
- [9] Y. Cai, E. F. Haratsch, O. Mutlu, and K. Mai, "Threshold Voltage Distribution in NAND Flash Memory: Characterization, Analysis, and Modeling," in *DATE*, 2013.
- [10] Y. Cai, Y. Luo, E. F. Haratsch, K. Mai, and O. Mutlu, "Data Retention in MLC NAND Flash Memory: Characterization, Optimization, and Recovery," in *HPCA*, 2015.
- [11] Y. Cai, O. Mutlu, E. F. Haratsch, and K. Mai, "Program Interference in MLC NAND Flash Memory: Characterization, Modeling, and Mitigation," in *ICCD*, 2013.
- [12] Y. Cai, G. Yalcin, O. Mutlu, E. F. Haratsch, A. Cristal, O. Unsal, and K. Mai, "Flash Correct and Refresh: Retention Aware Management for Increased Lifetime," in *ICCD*, 2012.
- [13] Y. Cai, G. Yalcin, O. Mutlu, E. F. Haratsch, A. Cristal, O. Unsal, and K. Mai, "Error Analysis and Retention-Aware Error Management for NAND Flash Memory," *Intel Technology Journal (ITJ)*, 2013.
- [14] Y. Cai, G. Yalcin, O. Mutlu, E. F. Haratsch, O. Unsal, A. Cristal, and K. Mai, "Neighbor Cell Assisted Error Correction in MLC NAND Flash Memories," in *SIGMETRICS*, 2014.
- [15] Y. Cai, S. Ghose, Y. Luo, K. Mai, O. Mutlu, and E. F. Haratsch, "Vulnerabilities in MLC NAND Flash Memory Programming: Experimental Analysis, Exploits, and Mitigation Techniques," in *HPCA*, 2017.
- [16] Y. Cai, Y. Luo, S. Ghose, and O. Mutlu, "Read Disturb Errors in MLC NAND Flash Memory: Characterization, Mitigation, and Recovery," in *DSN*, 2015.
- [17] B. Choi *et al.*, "Comprehensive Evaluation of Early Retention (Fast Charge Loss Within a Few Seconds) Characteristics in Tube-Type 3-D NAND Flash Memory," in *VLSIT*, 2016.
- [18] L. Chua, "Memristor—The Missing Circuit Element," *TCT*, 1971.
- [19] J. Cooke, "The Inconvenient Truths of NAND Flash Memory," *Flash Memory Summit*, 2007.
- [20] Y. Du, Q. Li, L. Shi, D. Zou, H. Jin, and C. J. Xue, "Reducing LDPC Soft Sensing Latency by Lightweight Data Refresh for Flash Read Performance Improvement," in *DAC*, 2017.
- [21] H. H. Frost, C. J. Camp, T. J. Fisher, J. A. Fuxa, and L. W. Shelton, "Efficient Reduction of Read Disturb Errors in NAND Flash Memory," U.S. Patent 7,818,525, 2010.
- [22] A. Fukami, S. Ghose, Y. Luo, Y. Cai, and O. Mutlu, "Improving the Reliability of Chip-Off Forensic Analysis of NAND Flash Memory Devices," *Digital Investigation*, vol. 20, pp. S1–S11, 2017.
- [23] T. G. Gary Tressler and P. Breen, "Read Disturb Characterization for Next-Generation Flash Systems," in *Flash Memory Summit*, 2015.

- [24] A. Goda and K. Parat, "Scaling Directions for 2D and 3D NAND Cells," in *IEDM*, 2012.
- [25] A. Grossi, C. Zambelli, and P. Olivo, "Reliability of 3D NAND Flash Memories," in *3D Flash Memories*. Springer, 2016, pp. 29–62.
- [26] L. M. Grupp, A. M. Caulfield, J. Coburn, S. Swanson, E. Yaakobi, P. H. Siegel, and J. K. Wolf, "Characterizing Flash Memory: Anomalies, Observations, and Applications," in *MICRO*, 2009.
- [27] D. Gruss, M. Lipp, M. Schwarz, D. Genkin, J. Juffinger, S. O'Connell, W. Schoechl, and Y. Yarom, "Another Flip in the Wall of Rowhammer Defenses," arXiv:1710.00551 [cs.CR], 2017.
- [28] D. Gruss, C. Maurice, and S. Mangard, "Rowhammer.js: A Remote Software-Induced Fault Attack in JavaScript," in *DIMVA*, 2016.
- [29] K. Ha, J. Jeong, and J. Kim, "A Read-Disturb Management Technique for High-Density NAND Flash Memory," in *APSys*, 2013.
- [30] K. Ha, J. Jeong, and J. Kim, "An Integrated Approach for Managing Read Disturbs in High-Density NAND Flash Memory," *TCAD*, vol. 35, no. 7, pp. 1079–1091, 2016.
- [31] H. Hassan, N. Vijaykumar, S. Khan, S. Ghose, K. Chang, G. Pekhimenko, D. Lee, O. Ergin, and O. Mutlu, "SoftMC: A Flexible and Practical Open-Source Infrastructure for Enabling Experimental DRAM Studies," in *HPCA*, 2017.
- [32] D. Ielmini, A. L. Lacaita, and D. Mantegazza, "Recovery and Drift Dynamics of Resistance and Threshold Voltages in Phase-Change Memories," *TED*, 2007.
- [33] J. Im *et al.*, "A 128Gb 3b/Cell V-NAND Flash Memory with 1Gb/s I/O Rate," in *ISSCC*, 2015.
- [34] Y. Jang, J. Lee, S. Lee, and T. Kim, "SGX-Bomb: Locking Down the Processor via Rowhammer Attack," in *SysTEX*, 2017.
- [35] L. Jiang, Y. Zhang, and J. Yang, "Mitigating Write Disturbance in Super-Dense Phase Change Memories," in *DSN*, 2014.
- [36] A. Jog, A. K. Mishra, C. Xu, Y. Xie, V. Narayanan, R. Iyer, and C. R. Das, "Cache Revive: Architecting Volatile STT-RAM Caches for Enhanced Performance in CMPs," in *DAC*, 2012.
- [37] D. Kang *et al.*, "7.1 256Gb 3b/cell V-NAND Flash Memory With 48 Stacked WL Layers," in *ISSCC*, 2016.
- [38] J. Katcher, "Postmark: A New File System Benchmark," Network Appliance, Tech. Rep. TR3022, 1997.
- [39] W.-S. Khwa *et al.*, "A Resistance-Drift Compensation Scheme to Reduce MLC PCM Raw BER by Over 100x for Storage-Class Memory Applications," in *ISSCC*, 2016.
- [40] N. Kim and J.-H. Jang, "Nonvolatile Memory Device, Method of Operating Nonvolatile Memory Device and Memory System Including Nonvolatile Memory Device," U.S. Patent 8,203,881, 2012.
- [41] Y. Kim, "Architectural Techniques to Enhance DRAM Scaling," Ph.D. dissertation, Carnegie Mellon Univ., 2015.
- [42] Y. Kim, R. Daly, J. Kim, C. Fallin, J. H. Lee, D. Lee, C. Wilkerson, K. Lai, and O. Mutlu, "Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors," in *ISCA*, 2014.
- [43] R. Koller and R. Rangaswami, "I/O Deduplication: Utilizing Content Similarity to Improve I/O Performance," *TOS*, 2010.
- [44] E. Kültürsay, M. Kandemir, A. Sivasubramaniam, and O. Mutlu, "Evaluating STT-RAM as an Energy-Efficient Main Memory Alternative," in *ISPASS*, 2013.
- [45] B. C. Lee, E. Ipek, O. Mutlu, and D. Burger, "Architecting Phase Change Memory as a Scalable DRAM Alternative," in *ISCA*, 2009.
- [46] B. C. Lee, E. Ipek, O. Mutlu, and D. Burger, "Phase Change Memory Architecture and the Quest for Scalability," *CACM*, 2010.
- [47] B. C. Lee, P. Zhou, J. Yang, Y. Zhang, B. Zhao, E. Ipek, O. Mutlu, and D. Burger, "Phase-Change Technology and the Future of Main Memory," *IEEE Micro*, 2010.
- [48] C.-Y. Liu, Y.-M. Chang, and Y.-H. Chang, "Read Leveling for Flash Storage Systems," in *SYSTOR*, 2015.
- [49] R.-S. Liu, C.-L. Yang, and W. Wu, "Optimizing NAND Flash-Based SSDs via Retention Relaxation," in *FAST*, 2012.
- [50] Y. Luo, Y. Cai, S. Ghose, J. Choi, and O. Mutlu, "WARM: Improving NAND Flash Memory Lifetime With Write-Hotness Aware Retention Management," in *MSST*, 2015.
- [51] Y. Luo, "Architectural Techniques for Improving NAND Flash Memory Reliability," Ph.D. dissertation, Carnegie Mellon Univ., 2018.
- [52] Y. Luo, S. Ghose, Y. Cai, E. F. Haratsch, and O. Mutlu, "Enabling Accurate and Practical Online Flash Channel Modeling for Modern MLC NAND Flash Memory," *JSAC*, 2016.
- [53] 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 *HPCA*, 2018.
- [54] C. Manning, "Yaffs NAND Flash Failure Mitigation," <http://www.yaffs.net/sites/yaffs.net/files/YaffsNandFailureMitigation.pdf>, 2012.
- [55] J. Meza, Y. Luo, S. Khan, J. Zhao, Y. Xie, and O. Mutlu, "A Case for Efficient Hardware/Software Cooperative Management of Storage and Memory," in *WEED*, 2013.
- [56] J. Meza, Q. Wu, S. Kumar, and O. Mutlu, "A Large-Scale Study of Flash Memory Failures In The Field," in *SIGMETRICS*, 2015.
- [57] R. Micheloni, Ed., *3D Flash Memories*. Dordrecht, Netherlands: Springer Netherlands, 2016.
- [58] R. Micheloni, S. Aritome, and L. Crippa, "Array Architectures for 3-D NAND Flash Memories," *Proc. IEEE*, Sep. 2017.
- [59] N. Mielke, T. Marquart, N.Wu, J.Kessenich, H. Belgal, E. Schares, and F. Triverdi, "Bit Error Rate in NAND Flash Memories," in *IRPS*, 2008.
- [60] K. Mizoguchi, T. Takahashi, S. Aritome, and K. Takeuchi, "Data-Retention Characteristics Comparison of 2D and 3D TLC NAND Flash Memories," in *IMW*, 2017.
- [61] O. Mutlu, "The RowHammer Problem and Other Issues We May Face as Memory Becomes Denser," in *DATE*, 2017.
- [62] O. Mutlu, "Memory Scaling: A Systems Architecture Perspective," in *IMW*, 2013.
- [63] O. Mutlu and L. Subramanian, "Research Problems and Opportunities in Memory Systems," *SUPERFRI*, 2014.
- [64] H. Naeimi, C. Augustine, A. Raychowdhury, S.-L. Lu, and J. Tschanz, "STT-RAM Scaling and Retention Failure," *Intel Technology Journal*, 2013.
- [65] D. Narayanan, A. Donnelly, and A. Rowstron, "Write off-Loading: Practical Power Management for Enterprise Storage," *TOS*, 2008.
- [66] I. Narayanan, D. Wang, M. Jeon, B. Sharma, L. Caulfield, A. Sivasubramaniam, B. Cutler, J. Liu, B. Khessib, and K. Vaid, "SSD Failures in Datacenters: What? When? and Why?" in *SYSTOR*, 2016.
- [67] 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 *IMW*, 2016.
- [68] N. Papandreou, T. Parnell, H. Pozidis, T. Mittelholzer, E. Eleftheriou, C. Camp, T. Griffin, G. Tressler, and A. Walls, "Using Adaptive Read Voltage Thresholds to Enhance the Reliability of MLC NAND Flash Memory Systems," in *GLSVLSI*, 2014.
- [69] K.-T. Park *et al.*, "A 7MB/s 64Gb 3-Bit/Cell DDR NAND Flash Memory in 20nm-Node Technology," in *ISSCC*, 2011.
- [70] K. Park *et al.*, "Three-Dimensional 128 Gb MLC Vertical NAND Flash Memory With 24-WL Stacked Layers and 50 MB/s High-Speed Programming," *J. Solid-State Circuits*, Jan. 2015.
- [71] T. Parnell and R. Pletka, "NAND Flash Basics & Error Characteristics – Why Do We Need Smart Controllers?" in *Flash Memory Summit*, 2017.
- [72] T. Parnell, N. Papandreou, T. Mittelholzer, and H. Pozidis, "Modelling of the Threshold Voltage Distributions of Sub-20nm NAND Flash Memory," in *GLOBECOM*, 2014.
- [73] A. Pirovano, A. L. Lacaita, F. Pellizzer, S. A. Kostylev, A. Benvenuti, and R. Bez, "Low-Field Amorphous State Resistance and Threshold Voltage Drift in Chalcogenide Materials," *TED*, 2004.
- [74] D. Poddebniak, J. Somorovsky, S. Schinzel, M. Lochter, and P. Rösler, "Attacking Deterministic Signature Schemes Using Fault Attacks," Cryptology ePrint Archive, Report 2017/1014, 2017.
- [75] M. K. Qureshi, V. Srinivasan, and J. A. Rivers, "Scalable High Performance Main Memory System Using Phase-Change Memory Technology," in *ISCA*, 2009.
- [76] K. Razavi, B. Gras, E. Bosman, B. Preneel, C. Guiffrida, and H. Bos, "Flip Feng Shui: Hammering a Needle in the Software Stack," in *USENIX Security*, 2016.
- [77] B. Schroeder, A. Merchant, and R. Lagisetty, "Reliability of NAND-Based SSDs: What Field Studies Tell Us," *Proc. IEEE*, 2017.
- [78] M. Seaborn and T. Dullien, "Exploiting the DRAM Rowhammer Bug to Gain Kernel Privileges," Google Project Zero Blog, 2015.
- [79] M. Seaborn and T. Dullien, "Exploiting the DRAM Rowhammer Bug to Gain Kernel Privileges," in *BlackHat*, 2015.
- [80] S. Sills, S. Yasuda, A. Calderoni, C. Cardon, J. Strand, K. Aratani, and N. Ramaswamy, "Challenges for High-Density 16Gb ReRAM with 27nm Technology," in *VLSIC*, 2015.
- [81] S. Sills, S. Yasuda, J. Strand, A. Calderoni, K. Aratani, A. Johnson, and N. Ramaswamy, "A Copper ReRAM Cell for Storage Class Memory Applications," in *VLSIT*, 2014.
- [82] C. W. Smullen, V. Mohan, A. Nigam, S. Gurumurthi, and M. R. Stan, "Relaxing Non-Volatility for Fast and Energy-Efficient STT-RAM Caches," in *HPCA*, 2011.
- [83] Storage Network Industry Assn., "IOTTA Repository: Cello 1999." <http://iotta.snia.org/traces/21>
- [84] D. B. Strukov, G. S. Snider, D. R. Stewart, and R. S. Williams, "The Missing Memristor Found," *Nature*, 2008.
- [85] T. Sugahara and T. Furuichi, "Memory Controller for Suppressing Read Disturb When Data Is Repeatedly Read Out," US Patent No. 8725952, 2014.
- [86] Z. Sun, X. Bi, H. Li, W.-F. Wong, Z.-L. Ong, X. Zhu, and W. Wu, "Multi Retention Level STT-RAM Cache Designs with a Dynamic Refresh Scheme," in *MICRO*, 2011.
- [87] K. Takeuchi, S. Satoh, T. Tanaka, K. Imamiya, and K. Sakui, "A Negative Vth Cell Architecture for Highly Scalable, Excellently Noise-Immune, and Highly Reliable NAND Flash Memories," *JSSC*, 1999.
- [88] Toshiba, "Toshiba Develops World's First 256Gb, 48-Layer BiCS FLASH," <http://toshiba.semicon-storage.com/us/company/taec/news/2015/08/memory-20150803-1.html>, 2015.
- [89] Univ. of Massachusetts, "Storage: UMass Trace Repository."
- [90] V. van der Veen, Y. Fratanonio, M. Lindorfer, D. Gruss, C. Maurice, G. Vigna, H. Bos, K. Razavi, and C. Guiffrida, "Drammer: Deterministic Rowhammer Attacks on Mobile Platforms," in *CCS*, 2016.
- [91] H.-S. P. Wong, H.-Y. Lee, S. Yu, Y.-S. Chen, Y. Wu, P.-S. Chen, B. Lee, F. T. Chen, and M.-J. Tsai, "Metal-Oxide RRAM," *Proc. IEEE*, 2012.

- [92] H.-S. P. Wong, S. Raoux, S. Kim, J. Liang, J. P. Reifenberg, B. Rajendran, M. Asheghi, and K. E. Goodson, "Phase Change Memory," *Proc. IEEE*, 2010.
- [93] Y. Xiao, X. Zhang, Y. Zhang, and R. Teodorescu, "One Bit Flips, One Cloud Flops: Cross-VM Row Hammer Attacks and Privilege Escalation," in *USENIX Security*, 2016.
- [94] Q. Xiong, F. Wu, Z. Lu, Y. Zhu, Y. Zhou, Y. Chu, C. Xie, and P. Huang, "Characterizing 3D Floating Gate NAND Flash," in *SIGMETRICS*, 2017.
- [95] H. Yoon, J. Meza, N. Muralimanohar, N. P. Jouppi, and O. Mutlu, "Efficient Data Mapping and Buffering Techniques for Multi-Level Cell Phase-Change Memories," *TACO*, 2014.
- [96] J. H. Yoon, "3D NAND Technology: Implications to Enterprise Storage Applications," in *Flash Memory Summit*, 2015.
- [97] C. Zambelli, P. Olivo, L. Crippa, A. Marelli, and R. Micheloni, "Uniform and Concentrated Read Disturb Effects in Mid-1X TLC NAND Flash Memories for Enterprise Solid State Drives," in *IRPS*, 2017.
- [98] Z. Zhang, W. Xiao, N. Park, and D. J. Lilja, "Memory Module-Level Testing and Error Behaviors for Phase Change Memory," in *ICCD*, 2012.
- [99] P. Zhou, B. Zhao, J. Yang, and Y. Zhang, "A Durable and Energy Efficient Main Memory Using Phase Change Memory Technology," in *ISCA*, 2009.
- [100] Y. Zhu, F. Wu, Q. Xiong, Z. Lu, and C. Xie, "ALARM: A Location-Aware Redistribution Method to Improve 3D FG NAND Flash Reliability," in *NAS*, 2017.