

Improving the Reliability of Chip-Off Forensic Analysis of NAND Flash Memory Devices

Aya Fukami
National Police Agency, Japan
Carnegie Mellon University

Saugata Ghose
Carnegie Mellon University

Yixin Luo
Carnegie Mellon University

Yu Cai
Carnegie Mellon University

Onur Mutlu
ETH Zürich
Carnegie Mellon University

ABSTRACT

Digital forensic investigators often need to extract data from a seized device that contains NAND flash memory. Many such devices are physically damaged, preventing investigators from using automated techniques to extract the data stored within the device. Instead, investigators turn to *chip-off analysis*, where they use a thermal-based procedure to physically remove the NAND flash memory chip from the device, and access the chip directly to extract the raw data stored on the chip.

We perform an analysis of the errors introduced into multi-level cell (MLC) NAND flash memory chips *after the device has been seized*. We make two major observations. First, between the time that a device is seized and the time digital forensic investigators perform data extraction, a large number of errors can be introduced as a result of charge leakage from the cells of the NAND flash memory (known as *data retention errors*). Second, when thermal-based chip removal is performed, the number of errors in the data stored within NAND flash memory can increase by two or more orders of magnitude, as the high temperature applied to the chip greatly accelerates charge leakage. We demonstrate that the chip-off analysis based forensic data recovery procedure is quite destructive, and can often render most of the data within NAND flash memory uncorrectable, and, thus, unrecoverable.

To mitigate the errors introduced during the forensic recovery process, we explore a new hardware-based approach. We exploit a fine-grained read reference voltage control mechanism implemented in modern NAND flash memory chips, called *read-retry*, which can compensate for the charge leakage that occurs due to (1) retention loss and (2) thermal-based chip removal. The read-retry mechanism successfully reduces the number of errors, such that the original data can be fully recovered in our tested chips as long as the chips were not heavily used prior to seizure. We conclude that

the read-retry mechanism should be adopted as part of the forensic data recovery process.

1 INTRODUCTION

NAND flash memory continues to increase in popularity as a storage medium for a wide range of devices, such as smartphones, thumb drives, and solid-state drives (SSDs). As a result, digital forensic investigators have been encountering significantly more NAND flash memory based devices than before during the course of criminal investigations. When an operational NAND flash memory based device is received for analysis, investigators can use *logical data extraction*, where data can be read out using an interface provided by the device vendor [2]. Commercial software-based forensic acquisition tools automate logical data extraction [2], and can yield sufficient data from the device. Unfortunately, a device received as part of an investigation may be physically damaged, or the device may not provide an interface for data acquisition, and as a result, its data may be inaccessible using the automated software-based approach. In these cases, digital forensic investigators must physically remove the NAND flash memory chip from the printed circuit board (PCB) inside the device [24]. Once the chip has been removed, investigators can perform low-level analysis on the chip, at which point the data that was originally on the chip can potentially be recovered. This analysis method is commonly referred to as *chip-off analysis*.

Previous research on forensic low-level analysis of NAND flash memory chips has focused on reverse engineering the techniques implemented by the original NAND flash memory controllers, in order to access the data residing on the chip. Breeuwsma et al. [7] established a thorough forensic data recovery procedure from a NAND flash memory chip, going from acquiring the physical image of the data on the NAND flash memory chip to reconstructing the file system used to store the data, by reverse engineering multiple

techniques that are implemented in NAND flash memory controllers. For relatively old devices, where the number of raw bit errors that occurred within the device are low, this data recovery procedure is often sufficient. However, as NAND flash memory has scaled down aggressively in process technology node size to enable higher storage capacity, the number of raw bit errors that occur in NAND flash memory has increased by several orders of magnitude, as demonstrated experimentally [9, 10, 12]. As a result, the procedure proposed by Breeuwsma et al. [7] often cannot adequately recover the data in modern NAND flash memory.

In order to ensure that data retrieved from NAND flash memory by an end user does *not* contain any errors despite the increasing occurrence of raw bit errors, modern NAND flash memory controllers employ sophisticated *error-correcting codes* (ECC), such as BCH codes [5, 27, 35] or LDPC codes [26, 56]. These codes can correct up to a fixed number of raw bit errors for every read operation. Likewise, to maintain the integrity of digital evidence extracted from a device that uses NAND flash memory as its storage medium, forensic investigators need to correct errors that appear in the raw data extracted from the device. Thus, it is essential for the forensic data recovery procedure to extract the ECC information stored within the chip and use this information to correct the errors. In addition to ECC, many modern flash controllers employ *data randomization* techniques, where data that is written into memory is scrambled by XORing the data with a reproducible pseudo-random number, to reduce the impact of data value dependence on reliability [19, 29]. van Zandwijk [51] provides a mathematical approach to reverse engineer *both* the ECC and data randomization algorithms from the raw data that is extracted from the NAND flash memory chip.

However, even when the ECC and data randomization algorithms are correctly identified and used to decode the raw data extracted from the NAND flash memory chip, digital forensic examiners may find that many chunks of data contain more errors than the ECC algorithm is able to correct. The ECC codeword contains only enough information to correct up to e bits within an n -bit data chunk (we refer to e as the *error correction capability*). If the data chunk contains more than e errors, the errors are *uncorrectable*, and the data cannot be successfully recovered. This compromises the integrity of the recovered data for forensic analysis. Our goal in this work is to develop a methodology for chip-off analysis that improves the likelihood of reliable data recovery beyond what is possible by using ECC.

To this end, this paper first investigates the impact of current forensic techniques on the number of raw bit errors that exist within a NAND flash memory chip that is removed from the PCB in a device under investigation. In order to perform

chip-off analysis, forensic investigators follow best practices used by electronics manufacturers for their *rework process*, where manufacturers remove and replace faulty components in their products [28]. This procedure uses hot air to heat the chip just enough to melt the solder that connects the chip to the PCB, which allows the safe removal of the chip. We refer to this technique as *thermal-based chip removal*. We demonstrate that, even though the temperature used during chip removal is the minimal temperature necessary for the solder to reach its melting point (usually more than 200°C), this temperature is still high enough to introduce a very large number of *new* raw bit errors into the chip. Unfortunately, since the ECC data stored on-chip has a fixed error correction capability, the newly-introduced errors are often capable of rendering much of the data on a NAND flash memory chip unrecoverable using traditional low-level analysis techniques, such as those described by Breeuwsma et al. [7] and van Zandwijk [51].

To mitigate the impact of the new errors introduced during chip removal, we examine exploiting the state-of-the-art dynamic mechanisms implemented within the NAND flash memory itself to help reduce the raw bit error rate. In particular, we study the *read-retry* mechanism [10]. In a NAND flash memory cell, data is stored in the form of the *threshold voltage* of the cell's floating-gate transistor (i.e., the voltage at which a transistor turns on). In order to read data from the NAND flash memory cell, a *read reference voltage* is applied to the transistor. If the read reference voltage is higher than the threshold voltage, the cell turns *on*; otherwise, the cell turns *off*. As several prior works have shown [9, 10, 12, 15, 21, 23], the threshold voltage of a floating-gate transistor can *shift over time*, due to continuous leakage of charge from the transistor. Since the standard read operation uses the *default read reference voltage*, it is unable to account for such a threshold voltage shift, and thus the read operation introduces an error (e.g., we read a bit value 1 even though the stored value was 0). The flash controller can mitigate these errors by dynamically adjusting the read reference voltage to compensate for the threshold voltage shift. This mechanism is known as *read-retry* [10, 15]. Modern NAND flash memory chips include several variants of read-retry, which use different techniques to adjust the read reference voltage.

We evaluate the suitability of the read-retry mechanism for forensic analysis. Our findings show that read-retry can significantly reduce the errors inside a NAND flash memory chip that are introduced over the course of chip-off analysis (by as much as 94.6%). The read-retry mechanism eliminates uncorrectable data chunks for NAND flash memory chips that have not been heavily used and thus improves the reliability of chip-off analysis based forensic data recovery. The read-retry based error mitigation techniques discussed in

this paper can be implemented when extracting a raw image directly from a NAND flash memory chip during forensic analysis using a universal flash chip reader, as described in Breeuwsma et al. [7].

This paper makes the following contributions:

- We experimentally demonstrate that thermal-based chip removal, a commonly-used technique during the digital forensic analysis of a NAND flash memory based device, can introduce a large number of errors into the raw data stored within a NAND flash memory chip, increasing the raw bit error rate by as much as $259\times$ (Section 6.2). The additional errors can often be enough to prevent investigators from successfully extracting correct data from the device, hampering digital forensic analysis.
- We experimentally evaluate how the read-retry mechanism, provided by NAND flash memory manufacturers in modern flash memory chips to control the read reference voltage in a fine-grained manner, can reduce the raw bit error rate of data read from NAND flash memory by as much as 94.6% (Section 6.3).
- Our evaluations show that by incorporating read-retry based error mitigation into the forensic data recovery procedure, we can mitigate the errors introduced during thermal-based chip removal and successfully read out the data stored within a NAND flash memory chip, unless the chip has been heavily used (Section 6.3).

2 BACKGROUND: MLC NAND FLASH MEMORY

We first provide necessary background on the design and operation of multi-level cell (MLC) NAND flash memory, a very common variant of NAND flash memory available in a wide number of consumer devices today. More detailed information on the operation of flash memory can be found in prior works [9, 12, 16, 17, 21, 25, 34].

2.1 Flash Memory Organization

NAND flash memory stores data within an array of *flash cells*. A cell consists of a single *floating-gate transistor*, where the floating gate of the transistor can store some amount of charge, as shown in Figure 1. The charge stored (i.e., *trapped*) within the floating gate determines the *threshold voltage* at which the transistor turns on. Oxide layers are placed above and below the floating gate to prevent the stored charge from leaking out of the floating gate. To program a flash cell to a specific threshold voltage, a high voltage is applied to the transistor’s *control gate*.




Figure 1: A flash memory cell, which consists of a floating-gate transistor.

The threshold voltage can be programmed to a voltage level within a fixed range. This fixed range is split up into multiple *voltage windows*, or *states*, where each state represents a certain bit value. Older NAND flash memory devices were made up of *single-level cells* (SLC), where each flash cell stores a one-bit value (i.e., a 0 or a 1). In order to provide higher storage density, NAND flash memory manufacturers now use *multi-level cell* (MLC) technology. A multi-level cell stores *two* bits of data within a single cell. To do this, the voltage range is split into *four* states (ER, P1, P2, and P3), with each state corresponding to one of the data values 00, 01, 10, or 11, as shown in Figure 2. Due to variation during programming, the threshold voltage of cells programmed to the same state is distributed across the voltage window for the state. This results in a threshold voltage distribution of flash cells across the voltage range, as we show in Figure 2.




Figure 2: Threshold voltage distribution of cells in MLC NAND flash.

A NAND flash memory chip contains thousands of *flash blocks*, which are two-dimensional arrays of flash cells. Figure 3 shows the internal organization of a block. Each block contains dozens of rows (i.e., *wordlines*) of flash cells. All of the cells on the same wordline are read and programmed as a single group. MLC NAND flash memory partitions the two bits of data in each cell across two separate *pages* (the unit of data programmed at a time). As Figure 2 shows, the two bits stored within a multi-level cell are referred to as the *least significant bit* (LSB) and the *most significant bit* (MSB). The LSBs of all cells on one wordline form the LSB page of that wordline (e.g., Page 1 of Wordline 1 in Figure 3), and the MSBs of these cells form the MSB page (e.g., Page 4 of Wordline 1). For 2y-nm (i.e., 20–24nm) NAND flash memory, a single page consists of between 4–16KB of data [37]. Within each column of flash cells in the block, the sources and drains of the cells’ transistors are connected in series to

form a *bitline*. The cells on a bitline share a common ground on one end, and are connected to a sense amplifier on the other end. Read, program, and erase operations to the block are managed by a *flash controller*.




Figure 3: Organization of a NAND flash block.

2.2 Programming and Erasing Data

To program data into a flash cell, the cell needs to be in the *erased* state (i.e., no charge should be stored within the floating gate of the cell). During a program operation, electrons are injected into the floating gate by applying a high positive voltage to the control gate (see Figure 1 for a diagram of the flash cell). NAND flash memory uses a procedure known as *incremental step-pulse programming*, or ISPP [48]. During ISPP, the high programming voltage is applied for a very short period, known as a *step-pulse*. ISPP then checks the current voltage of the cell. ISPP repeats the process of applying a step-pulse and checking the voltage until the cell reaches its target threshold voltage. If we want to overwrite the data that currently exists in a cell, we must first erase the data in the cell. Within NAND flash memory, an erase operation is performed at block granularity (i.e., an entire block of flash cells is erased at once).

Over time, as a cell is programmed and erased, the cell begins to *wear out*, reducing its ability to reliably store charge within the floating gate [3, 41]. As this wearout is a result of the number of times a cell is programmed and erased, we quantify the degree of wearout in *program/erase (P/E) cycles*, as done in many prior works [9–17].

2.3 Reading Data from NAND Flash

To read a page of data from a block, the flash controller applies a *read reference voltage* to the cell’s control gate. If the threshold voltage of a cell is lower than the read reference voltage, the cell switches *on*; otherwise, the cell switches *off*. The read reference voltage used to read a cell depends on which page is being read from the wordline. As shown in Figure 2, to determine the LSB of a cell, the controller applies a single read reference voltage, V_b . If the threshold voltage of the cell is lower than V_b , the cell is in either the ER state

or the P1 state, and holds an LSB of 1; otherwise, the cell is either in the P2 state or the P3 state, and holds an LSB of 0. To determine the MSB of a cell, the controller applies *two* read reference voltages, V_a and V_c . The two voltages allow the controller to determine if a cell is in the P1/P2 states, and holds an MSB of 0, or if the cell is in the ER/P3 states, and holds an MSB of 1.

Since multiple cells are tied together on a single bitline, we must ensure that the cells that are *not being read* pass through the data that is being output from the cell that we want to read. In order to achieve this, the flash controller applies a *pass-through voltage* to the control gate of each unread cell (V_{pass} in Figure 2). The pass-through voltage is higher than any threshold voltage that can be stored within a flash cell. This ensures that a cell that is *not* being read is always turned *on* during the read operation, allowing the data of the cell that we *do* want to read to successfully reach the sense amplifier.

2.4 Correcting Errors

Data stored within flash cells can often contain errors. An error is introduced when the threshold voltage of a cell shifts outside of the voltage window to which the cell was originally programmed. There are a number of sources of errors within flash memory [9, 12], such as retention loss [11, 15], cell wearout [9, 10, 32], cell-to-cell program interference [13, 14], and read disturb [16]. As flash cells scale down to smaller process technology nodes, the total amount of charge that each cell can store decreases, which, in turn, increases the susceptibility of the flash cells to errors [53].

In order to combat the errors contained within the cells (which we refer to as *raw bit errors*), NAND flash memory makes use of *error-correcting codes* (ECC) [9, 55]. When data is programmed to a flash page, an ECC codeword is also written, which contains enough redundancy to correct e bits out of the n -bit data [51]. We refer to e as the *error correction capability* of the codeword. When the flash page is subsequently read, the ECC codeword is sent alongside the data to the flash controller. Inside the controller, both the data and the ECC codeword are input to the ECC logic, which checks for errors using the implemented error correction algorithm. Based on the results of the algorithm, the controller fixes the erroneous bits in the data, if any, and returns the corrected data value. If the data read from the page contains no more than e raw bit errors, the controller successfully returns correct data to the end user. If the data read from the page contains more than e raw bit errors, full data correction is not possible, and the page data is said to be *corrupted* (i.e., it is uncorrectable) [33, 47]. A block that contains corrupted

data is marked by the flash controller as a *bad block* [36], and is no longer used for storing data.

3 NAND FLASH ERROR SOURCES DURING FORENSIC ANALYSIS

As discussed in Section 2.4, errors can occur when the threshold voltage of a flash cell shifts. The probability of occurrence of such shifts has increased due to continued device scaling, which allows manufacturers to increase the flash storage density [53]. The reliability of NAND flash memory has been widely researched (e.g., [3, 6, 9–17, 20, 23, 32, 33, 40, 42–47, 52]). For a detailed overview of MLC NAND reliability, we refer the reader to prior works [12, 54]. Cai et al. [9] investigated multiple error factors on MLC NAND flash memory, including program/erase errors, program interference errors, retention errors, and read errors. Digital forensic investigators should assume that the majority of NAND flash memory devices they receive for analysis already contain multiple errors at the time the device is obtained.

In this section, we focus on the two major sources of errors that can be *introduced* during digital forensic analysis, *in addition to* the preexisting errors within the NAND flash memory device. First, a device is often stored unused for several days or even weeks before investigators are able to examine its contents, due to issues such as transport time or lab backlog [18]. During this time, additional *retention errors* can occur, which we discuss in Section 3.1. Second, for devices where thermal-based chip removal is required, a high temperature must be steadily applied to the device. This high temperature rapidly accelerates the effect of retention errors, as we discuss in Section 3.2.

3.1 Retention Errors

Retention errors in a flash cell occur when the charge stored within the floating gate of cell transistor *leaks*. Due to the structure of the cell transistor (see Figure 1), where the floating gate and substrate are separated by an oxide layer, a small amount of charge *tunnels* through the oxide, causing the leakage. This trend accelerates as the P/E cycle count increases [9, 53], as repeated programming and erasing of a cell degrades the oxide layer, which in turn allows charge to tunnel through the oxide layer at an increasing rate [15]. This tunneling occurs whether or not a NAND flash memory device is powered up, causing retention errors to accumulate when the device is stored unused with or without power.

Imagine a hypothetical case where a device was seized at a crime scene, and the device is received for analysis at a digital forensics lab three weeks later. Let us assume that this device had been used over the course of three years, and that the NAND flash memory within the device endured

10 P/E cycles each day during these three years of usage, adding up to a total of approximately 10^4 P/E cycles over its lifetime. Figure 4 shows the relationship between the P/E cycle count and the retention error rate found by Cai et al. [9] for real 3x-nm MLC NAND flash memory chips. The figure demonstrates that the error rate grows with (1) the P/E cycle count and (2) the *retention age* (i.e., the time elapsed since the data was programmed). As we can see by comparing the curve labeled *1 Day* with the curve labeled *3 Weeks*, the number of errors in the device increases by 38× over the three-week period between device seizure and lab delivery, while the device is being stored and transported.




Figure 4: Raw bit error rate of a 3x-nm NAND flash memory chip for different retention ages vs. P/E cycle count. Reproduced from [9].

To provide more insight into retention errors, we perform an experimental analysis on the retention error rate of modern 2y-nm (i.e., 20–24nm) MLC NAND flash memory in Section 6.1.

3.2 Thermal Effect on Error Rate

After the device is received by the digital forensics lab, investigators must extract the data from the device. In many such cases, the device might have been damaged prior to seizure, and cannot be accessed using software-based analysis techniques. In these cases, a lower-level *chip-off analysis* [4] must be performed, where investigators physically remove the NAND flash memory chips from the device and use hardware that can extract data directly from the chips [24]. In order to remove the chips, investigators must melt the solder connecting the chips to the PCB, at which point the chips can be pulled off.

Unfortunately, this thermal-based chip removal procedure greatly accelerates the number of retention errors that occur. Mielke et al. [39] states that when NAND flash memory is exposed to high temperature, Arrhenius’ Law [1, 31] can be used to convert the effects of high temperature into additional data retention time at normal operating temperature

for the memory. Let t_b denote the amount of time that heat is applied to the chip, and t_r denote the equivalent retention age at the normal operating temperature. According to Arrhenius' Law, t_r can be calculated as:

$$t_r = t_b \cdot \exp\left[\frac{E_a}{k} \left(\frac{1}{T_r} - \frac{1}{T_b}\right)\right] \quad (1)$$

where k is the Boltzmann constant, T_r denotes the normal operation temperature, and T_b denotes the *baking temperature* (e.g., the high temperature applied to the chip during the removal process). E_a is the activation energy, which we set to 1.1 eV according to [38].

Following standard rework procedures used by electronics manufacturers [28], thermal-based chip removal requires investigators to apply 250°C of heat for a duration of approximately two minutes. Using Equation 1, we see that applying this heat adds the same number of retention errors as if we had left the NAND flash memory untouched at room temperature for 833 years. Given that the amount of retention errors after *only three years*, as shown in Figure 4, already introduces more than 1000× the number of errors into NAND flash memory, compared to the number of retention errors after one day, we can extrapolate that 833 years' worth of retention errors would be significantly and prohibitively larger (over 10⁵× the number of retention errors after one day). Such a large rate of errors could easily overwhelm the error correction capability of ECC algorithms employed in modern flash controllers.

We experimentally quantify the exact impact of high-temperature baking on modern 2y-nm MLC NAND flash memory in Section 6.2.

4 REDUCING NAND FLASH ERRORS WITH READ-RETRY

As we saw in Section 3, the number of errors that exist in the raw data can increase by several orders of magnitude even under best practices used in forensic investigation techniques. If left unaddressed, the number of errors can quickly exceed the total error correction capability of the flash device, which in turn results in partial or complete data loss during forensic data recovery. In this section, we examine a mechanism that exists in modern MLC NAND flash memory chips, called *read-retry*, which can be used as part of the recovery process to mitigate the high number of errors introduced during data recovery.

Recall from Section 2 that errors occur when the threshold voltage of a flash cell shifts, causing the read reference voltages to incorrectly interpret the state of the cell. Prior work has demonstrated that retention errors are the dominant source of errors in MLC NAND flash memory [9]. As

a result, the threshold voltage of a flash cell tends to *decrease* as the data retention age increases. To combat this decrease in threshold voltage, NAND flash memory manufacturers provide a *read-retry operation*, which can adjust the read reference voltages used to read data from a cell, and thus potentially reduce the number of raw bit errors in the data [10].

Figure 5 shows an example of a threshold voltage distribution that has shifted downwards due to charge leakage over retention time. For the sake of simplicity, we show only the ER state and P1 state distributions in the figure. If we had applied the normal read reference voltages (i.e., V_a , V_b , and V_c in Figure 2) to the original distribution (i.e., before the distribution shifted), we would be able to read out the values of all cells in the distribution correctly. Once the distribution shifts, the read reference voltages no longer fall in between the shifted distributions of each voltage state, but instead fall *inside* the distributions of some of the states. For example, we use the default read reference voltage V_a to distinguish between cells in the ER state and cells in the P1 state. We see in Figure 5 that, after the distribution shifts, some of the cells that belong to the P1 state now fall to the left of V_a . If the controller still uses V_a , it incorrectly classifies these cells as being in the ER state. As we can see, the default read reference voltages (V_a , V_b , V_c) can introduce many raw bit errors when the threshold voltage distribution shifts.




Figure 5: Effect of read-retry operation on a shifted threshold voltage distribution (showing the distributions for only the ER state and the P1 state).

We can apply the read-retry operation to our example shifted distribution to compensate for the effects of the voltage distribution shifts caused by charge leakage. The basic goal of the read-retry operation is to adjust the read reference voltages up or down *with the goal* of minimizing the errors *that are introduced* due to misclassification. We denote the *optimal read reference voltages* (i.e., the voltages that are exactly in the middle of the distance between two neighboring distributions, which minimizes the number of errors) as V_a^{opt} , V_b^{opt} , and V_c^{opt} . One example read-retry mechanism can adjust the voltages down one step at a time during a read operation, checking to see whether the number of errors goes down with each subsequent step. We illustrate how this

example mechanism works in Figure 5. Here, the read-retry mechanism tries to adjust the voltage used to distinguish between cells in the ER state and cells in the P1 state. The mechanism tries several voltages (V_{RRn} in the figure, where n represents the n th voltage tried). As we can see in Figure 5, the mechanism eventually finds a voltage (V_{RR3}) where there are no cell classification errors (because V_{RR3} falls *between* the threshold voltage distributions of the ER and P1 states). Even though V_{RR3} is higher than V_a^{opt} , we can still use V_{RR3} to safely extract data from the NAND flash memory *without any errors*.

Note that the read-retry operation itself does *not always* reduce errors. While retention errors usually shift the threshold voltage of a cell down, the threshold voltage of a cell can sometimes *increase* from its original state, as a result of cell-to-cell interference that occurs during programming, reading, and erasing [13, 14, 16]. The direction and magnitude of the change in threshold voltage vary for each cell, due to factors such as the retention age of the data in the cell, the number of read and program operations performed to neighboring cells, and manufacturing process variation. As a result, the threshold voltage distributions of each state do *not* shift uniformly, and can overlap with each other. Therefore, it is possible that simply shifting the read reference voltage up or down (without checking the resulting number of errors) could unintentionally introduce more errors than the normal read that is done with the default read reference voltage. A read-retry based mechanism that uses the flash controller to check the change in the number of errors for every tried read reference voltage can ensure that it never increases the number of errors during a read operation.

Some NAND flash memory chips expose a number of different modes for the read-retry operation. The details of these modes are, unfortunately, often publicly unavailable, making it difficult to know, with certainty, whether a mode supports the ability to ensure that the number of read errors never increases. As we show in Section 6, since the thermal-based chip removal procedure introduces a large number of retention errors (i.e., it shifts the threshold voltage distribution down significantly), we can improve the read error rate even for read-retry modes that do *not* check or expose the error rate change after a read is performed (though the improvements could be larger if the read-retry mode exposes more information to the controller).

5 TESTING METHODOLOGY

In order to investigate the raw bit errors that digital forensic investigators would encounter during their chip-off analysis (see Section 3.2), and to evaluate the effectiveness of using the read-retry operation (see Section 4) to overcome these errors,

we examine the effects of data retention and thermal-based chip removal using two new 2y-nm NAND flash memory chips manufactured by two different vendors. We refer to the two chips in the rest of the paper as Chip A and Chip B. We use an Altera DE0 field-programmable gate array (FPGA) board [49] to design a controller that communicates with the target chips, similar to prior work [7, 8]. Figure 6 shows a photograph of our testing environment. The FPGA issues commands to and receives data from the target NAND flash memory chip. The USB microcontroller sends the data received from the NAND flash memory chip to the PC, where we collect and analyze the data. We conduct all of our testing at room temperature, unless otherwise noted.




Figure 6: Photograph of our test infrastructure used to extract data from MLC NAND flash memory chips.

First, we examine the retention errors that forensic examiners would encounter when a device is received for forensic examination. We choose multiple blocks randomly from the target chips, and we program pseudo-random data into each block. We use pseudo-random data to mimic the data scrambling employed by modern flash controllers [19, 29]. To simulate the different impact that retention errors have on relatively new devices and on heavily-used devices, we divide the blocks into multiple subsets, and perform a fixed and different number of P/E operations/cycles to each subset. Each program operation writes pseudo-random data to the blocks. We choose P/E cycle counts of 10, 300, 1000, 2500, and 4000 to examine a wide range of device wear. Once all of the blocks reach their target P/E cycle count, we perform reads at set retention ages to capture the effects of retention age on the error rate. We study six retention ages: Day 0 (i.e., immediately after programming), Day 1, and Weeks 1, 2, 3, and 4.

Next, we simulate the thermal-based chip removal procedure by baking the chips once the data stored in the chips reaches a certain retention age. Using the same parameters that are used for chip removal during chip-off analysis, we apply a temperature of 250°C to the chips for two minutes, using a heat gun. After this baking procedure is complete, we immediately perform one more set of data read operations

to measure the error rate. We study the effects of baking at three different retention ages: Day 0 (i.e., immediately after programming), Week 1, and Week 4.

One of the target chips (Chip B) has the read-retry operation implemented. For this chip, we perform multiple read operations at every retention age, performing one read for each available read-retry mode, and one read without read-retry. The read-retry modes in Chip B do not enable us to check or observe whether the number of errors is lower than that when the default read reference voltages are used. Each mode simply shifts the read reference voltage to a different fixed level.

6 EXPERIMENTAL RESULTS

We evaluate the reliability impact of chip-off digital forensic analysis on real MLC NAND flash memory chips. First, we examine the effects of data retention on the raw bit error rate (RBER) of two real MLC NAND flash memory chips (Section 6.1). Next, we examine how applying high temperature to the two chips, under the same conditions as thermal-based chip removal in chip-off analysis, affects the raw bit error rate (Section 6.2). Finally, we demonstrate how the error rate changes when we use the read-retry mechanism available on one of the chips (Section 6.3).

Note that error characteristics differ by manufacturer and by chip model, and thus the RBER of *other* NAND flash memory chips may *not* always be consistent with the results that we observe here. However, given that digital devices are exposed to varying levels of stress (e.g., mobile phones left in a car on a hot day, solid-state drives salvaged from CCTV cameras found at a fire scene, dashcams that write to and erase data from NAND flash memory constantly), digital forensic investigators should assume that the NAND flash memory they are investigating can possibly contain *even more errors* than what we show in this paper.

6.1 Errors Introduced Due to Retention Time

Figure 7 shows how the RBER of the tested NAND flash memory chips varies with (1) the P/E cycle count of the chip (i.e., how much the flash memory cells on the chip have been worn out), and (2) the *retention age* of the data (i.e., the amount of time that elapses after the data was written). We make three major observations about the impact of retention age on MLC NAND flash memory.

First, we observe that in both chips, the RBER grows as the retention age increases. The Day 0 results in Figure 7 show the number of errors that exist in flash blocks immediately after the data is programmed to the NAND flash chips. We see that the RBER increases over time, but that the largest

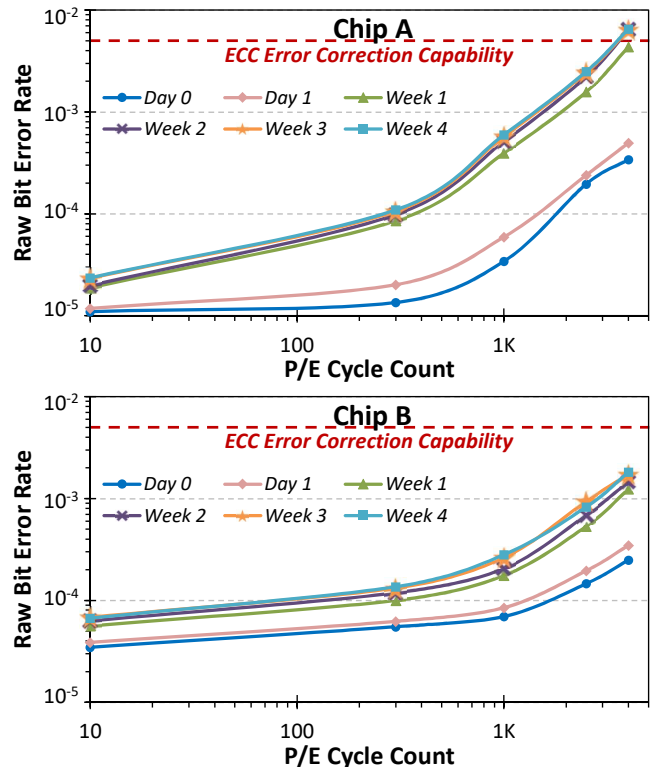


Figure 7: Raw bit error rate at different data retention ages, for different P/E cycle counts.

increases occur soon after the data is programmed. For example, the increase in RBER is much greater during the first week (6.34 \times for Chip A and 1.81 \times for Chip B at 300 P/E cycles, compared to Day 0 of each respective chip) than the second week (1.15 \times for Chip A and 1.19 \times for Chip B), and greater during the second week than the third week (1.08 \times for Chip A and 1.12 \times for Chip B).

Second, we observe that in both chips, the RBER grows as the P/E cycle count increases. In other words, as a NAND flash memory chip is worn out, its susceptibility to raw bit errors due to retention increases, for data with the same retention age. Note that the y-axis in Figure 7 is in log scale. A chip at a higher P/E cycle count (i.e., a chip with greater wearout) accumulates retention errors at a much faster rate than a chip at a lower P/E cycle count [15].

Third, we observe that while the RBER grows with both wearout and retention age, the overall RBER of the chip does *not* exceed the error correction capability of ECC unless the P/E cycle count significantly exceeds the endurance guaranteed by manufacturers. For 2y-nm MLC NAND flash memory, a controller that employs BCH codewords [5, 27, 35] for ECC can typically correct 40 bits of errors for every 1KB of data [22] (i.e., it can correct errors for an RBER of up

to 4.9×10^{-3}). As our results from Figure 7 show, the overall RBER stays lower than this error correction capability through a retention age of four weeks, for P/E cycle counts below 3000 cycles, which is the typical endurance of commercial 2y-nm MLC NAND flash memory [22].

Uncorrectable Data. We implement a 70-byte BCH codeword for ECC within our experimental platform, to provide the expected 40 bits of correction capability for each 1KB chunk of data. Note that a data chunk is smaller than a page [51], which can range from 4–16KB of data [37]. When we read data from our test chips, we use the BCH codeword to determine how many of the data chunks *cannot* be successfully corrected by ECC.

Table 1 shows the fraction of pages that contain at least one uncorrectable data chunk. As shown in Table 1, even before chip-off analysis is performed, if a chip has been worn out significantly (e.g., after 2500 P/E cycles for Chip A), it can contain some uncorrectable pages even at a retention age of just one week. However, for less worn-out chips (e.g., a chip at 1000 P/E cycles), none of the pages are uncorrectable even after a retention age of four weeks.

Table 1: Fraction of pages containing uncorrectable 1KB data chunks. A dash (–) indicates that no data chunks are uncorrectable.

Chip	Retention Age (days)	P/E Cycle Count				
		10	300	1000	2500	4000
A	0	–	–	–	–	–
	1	–	–	–	–	–
	7	–	–	–	0.9%	64.9%
	14	–	–	–	9.6%	92.2%
	21	–	–	–	14.4%	91.8%
	28	–	–	–	15.9%	91.9%
B	0	–	–	–	–	–
	1	–	–	–	–	–
	7	–	–	–	–	–
	14	–	–	–	–	–
	21	–	–	–	–	0.0039%
	28	–	–	–	–	0.0065%

For perspective, a device can reach 2190 P/E cycles if all of the pages in the NAND flash memory chip are written to twice a day, every day, over a period of three years. We observe that once we exceed the expected endurance of 2y-nm MLC NAND flash memory (which is 3000 P/E cycles), the fraction of pages with uncorrectable errors grows rapidly for Chip A. At 4000 P/E cycles, 64.9% of the pages in Chip A contain uncorrectable errors after a retention age of *only* one week.

We observe that even when the *overall* RBER stays lower than the ECC error correction capability, errors in *some* pages

become uncorrectable over time. For example, Chip A’s RBER at a retention age of one week (i.e., seven days) at 2500 P/E cycles is 1.6×10^{-3} (see Figure 7), which is lower than the ECC error correction capability of 4.9×10^{-3} . At the same time, 0.9% of the pages in Chip A contain uncorrectable errors, as we see in Table 1. We see similar behavior for Chip B, even though its overall RBER *always* remains below the error correction capability (see Figure 7).

We conclude that even if we assume that all of the data inside the device is refreshed immediately before the device was confiscated, a worn-out device can quickly accumulate errors, and some of those errors become uncorrectable over time (as we have shown in Figure 7 and Table 1). Thus, in order to avoid data loss due to uncorrectable data pages, data needs to be extracted from a NAND flash memory based device at the *earliest possible time after the receipt of the device*.

Error Characterization. By investigating the state of each cell at various retention ages, we can characterize a number of trends in the threshold voltage distribution shift. We say that a cell belongs to the set $[S_O, S_M]$ if the cell was originally programmed to state S_O , but is misread as belonging to state S_M when we use the default read reference voltages (see Section 4). The graphs in Figure 8 show the fraction of cells in Chip A that are in the set $[S_O, S_M]$, for all neighboring (S_O, S_M) pairs, out of the total number of cells originally programmed to state S_O , across a range of retention ages.




Figure 8: Fraction of cells in Chip A that were programmed to state S_O but are misread as belonging to state S_M , out of the total number of cells originally programmed to state S_O .

We observe from Figure 8 that when S_M is a *lower* voltage state than S_O , a greater number of cells belong to $[S_O, S_M]$ as

the retention age increases. For example, after a retention age of four weeks (i.e., 28 days), 0.20% of cells that were originally programmed to the P2 state are misread as belonging to the P1 state (i.e., the cells are in the set [P2, P1]), as opposed to only 0.02% after a retention age of one day. We find that regardless of the retention age, when S_M is a *higher* voltage state than S_O , only a very small number of cells belong to the set [S_O, S_M] (e.g., [ER, P1]). From these results, we find that the threshold voltage of a misread cell tends to be lower as the retention age increases. We conclude that the threshold voltage reduction, which occurs as a result of charge leakage from the floating gate of a flash memory cell, is the dominant source of errors that are introduced by retention age.

6.2 Errors Due to Thermal-Based Chip Removal

We now study how the RBER due to retention errors changes *after* the thermal-based chip removal procedure is performed as a part of chip-off analysis. Figure 9 shows the RBER after we perform the chip baking process (see Section 5) to emulate the removal procedure. Note that the RBER data for Day 0 (i.e., immediately after programming), Week 1, and Week 4 *before baking* is the same as the data shown in Figure 7.

We observe that simply applying the heat required for chip removal causes the RBER to increase significantly. When the heat is applied immediately after the data is written (*Day 0, After Baking* in the figure), at 1000 P/E cycles, the RBER increases by 432× for Chip A, and by 17× for Chip B, compared to the RBER before baking (*Day 0, Before Baking*). When the heat is applied four weeks after the data is written at the same P/E cycle count, (*Week 4, After Baking*), the RBER increases by 47× and 54× for Chips A and B, respectively, compared to the RBER before baking (*Week 4, Before Baking*). Starting at 300 P/E cycles, the RBER exceeds the ECC error correction capability when heat is applied to the chip.

The impact of the heat applied during chip removal can cause critical damage to the data stored within the NAND flash memory chip that is being analyzed. Suppose that a digital forensic investigator starts the chip-off procedure four weeks after a device has been seized, and that the NAND flash memory chip inside the device was only lightly used (e.g., 300 P/E cycles) prior to seizure. Before applying heat, the RBER remains safely within the error correction capability of contemporary ECC, as shown in Figure 9 (*Week 4, Before Baking* at 300 P/E cycles), with a raw bit error rate of 1.1×10^{-4} for Chip A and 1.4×10^{-4} for Chip B. However, after applying the chip removal temperature, the RBER exceeds the error correction capability of ECC (*Week 4, After Baking* at 300 P/E cycles). The RBER becomes 5.6×10^{-3} and 5.0×10^{-3} for Chip A and Chip B, respectively. At such a high RBER, it is




Figure 9: Raw bit error rate before and after baking NAND flash memory chips.

impossible to correct all of the errors in the data with the given ECC error correction capability. Thus, the integrity of the data recovery is compromised. As a point of comparison, we extrapolate the increase in RBER between Week 1 and Week 4 before baking for both chips, to see how long it would take for the chips to reach the RBER after baking if we had not baked the chips. Baking can increase the RBER by 113× for Chip A, and by 38× for Chip B on average, while the increase in RBER between *Week 1, Before Baking* and *Week 4, Before Baking* is 1.43× for both chips. Therefore, applying heat to a chip induces approximately *two to five years'* worth of retention errors at room temperature.

Table 2 shows the percentage of uncorrectable data chunks *after* the chips are baked (compare this to Table 1, which shows the uncorrectable data chunks *before* baking). Nearly all of the pages stored within the NAND flash memory contain uncorrectable errors after the baking process. At only 300 P/E cycles, 84.2% of the pages in Chip A and 83.6% of the pages in Chip B contain uncorrectable errors, and *all* pages contain errors when we reach 2500 P/E cycles for both Chip A and Chip B, when the heat is applied four weeks after the data is written. From our analysis, we conclude that, when left unmitigated, the thermal-based chip removal

procedure is prohibitively destructive, as it greatly decreases the amount of data that can be successfully retrieved from NAND flash memory during forensic recovery.

Table 2: Fraction of pages containing uncorrectable 1KB data chunks after applying heat to the chips. A dash (–) indicates that no data chunks are uncorrectable.

Chip	Retention Age (days)	P/E Cycle Count				
		10	300	1000	2500	4000
A	7	–	29.1%	99.8%	100.0%	100.0%
	28	0.7%	84.2%	96.9%	100.0%	100.00%
B	7	–	78.1%	96.5%	96.9%	96.9%
	28	–	83.6%	99.7%	100.0%	100.0%

6.3 Read-Retry Operation

We now investigate the ability of the read-retry operation (see Section 4) to mitigate the errors introduced during the thermal-based chip removal process. We are able to exploit the read-retry mechanism built into one of our chips (Chip B). In Chip B, we have access to two read-retry modes, which we refer to as Mode A and Mode B.¹

Figure 10 shows how the read-retry mechanism affects the RBER as the retention age increases, *ignoring the effects of the chip removal process for now*. We observe that across all P/E cycle counts, while the RBER increases with retention age when we use the default read operation, the RBER actually *decreases* with retention age if we use either of the read-retry modes. However, we find that the read-retry modes appear to be basic: *only* Mode A can outperform the default read operation *only* at high P/E cycle counts. We find that we do not benefit from read-retry for most of the time, as we are unable to check or control the behavior of the implemented read-retry modes when they lead to a higher RBER than the default read reference voltages.² Fortunately, without the thermal removal process, the RBER after using Mode A typically remains within the ECC error correction capability, as shown in Figure 7. When the chip is worn out beyond the manufacturer’s endurance specification (e.g., at 4000 P/E cycles), Mode A effectively reduces the RBER compared to the default read operation after a retention age of two weeks (i.e., 14 days). In fact, we find that some of the uncorrectable pages observed in Section 6.1 become correctable when we use Mode A (as we explain below).

¹No documentation is available from the manufacturer on how the modes operate.

²As explained in Section 4, it is difficult to know, with certainty, whether the read-retry modes implemented in the test chip check the error rate after read-retry is performed.

When we apply our baking process to emulate thermal-based chip removal, we find that the read-retry modes, even with their uncontrollable behavior, are very successful at reducing the RBER, as shown in Figure 11. As we saw in Section 6.2, baking a chip induces 2–5 years’ worth of retention errors. These errors are the result of a significant downward shift in the threshold voltage distribution of the flash cells. The two read-retry modes are able to adapt the read reference voltages to the shifted distribution such that *they can reduce the post-baking RBER to a level that is within the error correction capability of the implemented ECC algorithm*, even at high P/E cycle counts. For example, while the RBER for Chip B at 1000 P/E cycles (1.5×10^{-2}) is significantly over the ECC error correction capability after the chip is baked, Modes A and B can reduce the RBER by 88.6% (i.e., to 1.7×10^{-3}) and 94.6% (i.e., to 8.2×10^{-4}), respectively.

The read-retry mechanism significantly reduces the number of uncorrectable data chunks after the chip is baked. Table 3 shows the number of uncorrectable data chunks when we use the default read operation, and when we use the available read-retry modes, for data with a retention age of four weeks. We make two observations from the table. First, at low P/E cycle counts (e.g., 1000 P/E cycles), Mode B can *completely eliminate* the uncorrectable data chunks that were introduced by baking. Second, at high P/E cycle counts, while read-retry cannot fully eliminate the uncorrectable data chunks, it can significantly reduce the number of them. For example, at 2500 P/E cycles, the default read operation produces uncorrectable errors in *every* data chunk. Read-retry Mode B reduces the number of uncorrectable data chunks to 49.5% of all chunks. We believe these results can be significantly improved if the read-retry mode implemented by the chip is modified to expose more information to the flash controller.

Table 3: Fraction of pages containing uncorrectable 1KB data chunks after Chip B is baked, with and without read-retry. Values in bold indicate cases where read-retry eliminates all of the uncorrectable chunks.

Read Mode	P/E Cycle Count				
	10	300	1000	2500	4000
Default Read	0.0%	83.6%	99.7%	100.0%	100.0%
Read-Retry Mode A	0.0%	0.0%	12.1%	69.0%	99.1%
Read-Retry Mode B	0.0%	0.0%	0.0%	49.5%	90.6%

We conclude that digital forensic investigators should employ the read-retry mechanisms built into NAND flash memory chips, as read retry is able to overcome the large number of uncorrectable errors introduced due to the exposure of chips to very high temperatures during chip-off analysis.




Figure 10: Effect of read-retry modes on the RBER of Chip B as retention age increases.




Figure 11: RBER with read-retry after Chip B is baked.

7 RELATED WORK

To our knowledge, this is the first work to (1) demonstrate that the high temperatures used during thermal-based chip removal in chip-off analysis greatly increase the error rate in NAND flash memory chips, and is thus detrimental to forensic data recovery if left unmitigated; and (2) analyze the effect of the read-retry mechanism, commonly implemented in modern NAND flash memory chips, on recovering the data stored in the chips from a digital forensics perspective.

There are several works that characterize (1) various types of errors in MLC NAND flash memory (e.g., [9–16, 32, 42–46]), and (2) the effect of temperature on NAND flash memory reliability [50, 52]. None of these works (1) quantify the largely detrimental effects of high temperatures on reliable recovery of data from MLC NAND flash memory chips, (2) examine these issues from the perspective of digital forensic data recovery, or (3) demonstrate the value of read-retry in mitigating the NAND flash errors introduced as a result of thermal effects.

Other works in digital forensics analyze data recovery from flash memory chips [7, 30, 51]. However, none of these

examine the effects of (1) the powerful read-retry mechanism, or (2) high temperatures on error rates during the data recovery procedure.

8 CONCLUSION

With the increasing popularity of NAND flash memory as a storage medium, digital forensic investigators today are required to perform data recovery from a growing number of NAND flash memory based devices that are seized during the course of criminal investigations. Prior works have documented a series of procedures that can be used by investigators to extract data from physically-damaged devices, including the use of chip-off analysis. In this work, we find that the large amount of time that may elapse between device seizure and data extraction increases the error rate of data stored within the device. Often, this error rate can exceed the number of errors that can be corrected by the internal error correction mechanisms originally implemented by the NAND flash memory controller. In many cases, a majority of the data extracted from the device can contain uncorrectable errors. Therefore, we conclude that it is critical for digital forensic investigators to perform data extraction from NAND flash memory based digital devices at the earliest point of time after seizure of the target device.

In situations where chip-off analysis is required, a chip is exposed to high temperatures during the thermal-based chip removal process needed for chip-off analysis. We find that these high temperatures can further increase the error rate by more than two orders of magnitude, despite the use of best practices taken from electronics rework procedures. We demonstrate that using the read-retry mechanism built into modern NAND flash memory chips, instead of simply using the default read operation when extracting data, provides a promising solution. As our experimental results show, the read-retry mechanism can significantly mitigate the error rate increase caused by the thermal-based chip removal process. We conclude that forensic data recovery from NAND

flash memory should adopt the read-retry mechanism as part of the data recovery procedure.

REFERENCES

- [1] S. Arrhenius, "Über die Reaktionsgeschwindigkeit bei der Inversion von Rohrzucker durch Säuren," *Zeitschrift für Physikalische Chemie*, 1889.
- [2] R. Ayers *et al.*, *Guidelines on Mobile Device Forensics*, National Institute of Standards and Technology, 2014.
- [3] H. P. Belgal *et al.*, "A New Reliability Model for Post-Cycling Charge Retention of Flash Memories," in *IRPS*, 2002.
- [4] D. Billard and P. Vidonne, "Chip-off by matter subtraction: Frigida via," in *SADFE*, 2015.
- [5] R. Bose and D. Ray-Chaudhuri, "On a Class of Error Correcting Binary Group Codes," *Information and Control*, 1960.
- [6] A. Brand *et al.*, "Novel Read Disturb Failure Mechanism Induced by FLASH Cycling," in *IRPS*, 1993.
- [7] M. Breeuwsma *et al.*, "Forensics Data Recovery from Flash Memory," *Small Scale Digital Device Forensics Journal*, 2007.
- [8] Y. Cai *et al.*, "FPGA-Based Solid-State Drive Prototyping Platform," in *FCCM*, 2011.
- [9] Y. Cai *et al.*, "Error Patterns in MLC NAND Flash Memory: Measurement, Characterization, and Analysis," in *DATE*, 2012.
- [10] Y. Cai *et al.*, "Threshold Voltage Distribution in MLC NAND Flash Memory: Characterization, Analysis, and Modeling," in *DATE*, 2013.
- [11] Y. Cai *et al.*, "Flash Correct-and-Refresh: Retention-Aware Error Management for Increased Flash Memory Lifetime," in *ICCD*, 2012.
- [12] Y. Cai *et al.*, "Error Analysis and Retention-Aware Error Management for NAND Flash Memory," *Intel Technology Journal*, 2013.
- [13] Y. Cai *et al.*, "Program Interference in MLC NAND Flash Memory: Characterization, Modeling, and Mitigation," in *ICCD*, 2013.
- [14] Y. Cai *et al.*, "Neighbor-Cell Assisted Error Correction for MLC NAND Flash Memories," in *SIGMETRICS*, 2014.
- [15] Y. Cai *et al.*, "Data Retention in MLC NAND Flash Memory: Characterization, Optimization, and Recovery," in *HPCA*, 2015.
- [16] Y. Cai *et al.*, "Read Disturb Errors in MLC NAND Flash Memory: Characterization and Mitigation," in *DSN*, 2015.
- [17] Y. Cai *et al.*, "Vulnerabilities in MLC NAND Flash Memory Programming: Experimental Analysis, Exploits, and Mitigation Techniques," in *HPCA*, 2017.
- [18] E. Casey *et al.*, "Investigation Delayed Is Justice Denied: Proposals for Expediting Forensic Examinations of Digital Evidence," *Journal of Forensic Sciences*, 2009.
- [19] J. Cha and S. Kang, "Data Randomization Scheme for Endurance Enhancement and Interference Mitigation of Multilevel Flash Memory Devices," *ETRI Journal*, 2013.
- [20] A. Chimenton *et al.*, "Reliability of Floating Gate Memories," in *Error Correction Codes for Non-Volatile Memories*. Springer Netherlands, 2008.
- [21] L. Crippa and R. Micheloni, "MLC Storage," in *Inside NAND Flash Memories*. Springer Netherlands, 2010.
- [22] Cypress Semiconductor Corp., "SLC Versus MLC NAND Flash Memory," 2013. [Online]. Available: <http://www.cypress.com/file/209181/download>
- [23] R. Degraeve *et al.*, "Analytical Percolation Model for Predicting Anomalous Charge Loss in Flash Memories," *IEEE Transactions on Electron Devices*, 2004.
- [24] ENFSI, "Best Practice Manual for the Forensic Examination of Digital Technology," 2015. [Online]. Available: http://enfsi.eu/wp-content/uploads/2016/09/enfsi-bpm-fit-01_1.pdf
- [25] C. Friederich, "Program and Erase of NAND Memory Arrays," in *Inside NAND Flash Memories*. Springer Netherlands, 2010.
- [26] R. G. Gallager, *Low-Density Parity-Check Codes*. MIT Press, 1963.
- [27] A. Hocquenghem, "Codes Correcteurs d'Erreurs," *Chiffres*, 1959.
- [28] JEDEC Solid State Technology Association, "Moisture/Reflow Sensitivity Classification for Nonhermetic Solid State Surface Mount Devices," IPC/JEDEC J-STD-020D.1, 2007.
- [29] C. Kim *et al.*, "A 21 nm High Performance 64 Gb MLC NAND Flash Memory with 400 MB/s Asynchronous Toggle DDR Interface," *IEEE Journal of Solid-State Circuits*, 2012.
- [30] C. Klaver, "Windows Mobile Advanced Forensics," *Digital Investigation*, 2010.
- [31] K. J. Laidler, "The Development of the Arrhenius Equation," *Journal of Chemical Education*, 1984.
- [32] Y. Luo *et al.*, "Enabling Accurate and Practical Online Flash Channel Modeling for Modern MLC NAND Flash Memory," *IEEE Journal on Selected Areas in Communications*, 2016.
- [33] J. Meza *et al.*, "A Large-Scale Study of Flash Memory Failures in the Field," in *SIGMETRICS*, 2015.
- [34] R. Micheloni *et al.*, "NAND Overview: From Memory to Systems," in *Inside NAND Flash Memories*. Springer Netherlands, 2010.
- [35] R. Micheloni *et al.*, "BCH Hardware Implementation in NAND Flash Memories," in *Error Correction Codes for Non-Volatile Memories*. Springer, 2008.
- [36] Micron Technology, Inc., "Bad Block Management in NAND Flash Memory," 2011. [Online]. Available: https://www.micron.com/~media/documents/products/technical-note/nand-flash/tn2959_bbm_in_nand_flash.pdf
- [37] Micron Technology, Inc., "How Micron SSDs Handle Unexpected Power Loss," 2015. [Online]. Available: https://www.micron.com/~media/documents/products/white-paper/ssd_power_loss_protection_white_paper_lo.pdf
- [38] N. Mielke *et al.*, "Flash EEPROM Threshold Instabilities Due to Charge Trapping During Program/Erase Cycling," *IEEE Transactions on Device and Materials Reliability*, 2004.
- [39] N. Mielke *et al.*, "Recovery Effects in the Distributed Cycling of Flash Memories," in *IRPS*, 2006.
- [40] N. Mielke *et al.*, "Bit Error Rate in NAND Flash Memories," in *IRPS*, 2008.
- [41] Y. Pan *et al.*, "Exploiting Memory Device Wear-Out Dynamics to Improve NAND Flash Memory System Performance," in *FAST*, 2011.
- [42] N. Papandreou *et al.*, "Using Adaptive Read Voltage Thresholds to Enhance the Reliability of MLC NAND Flash Memory Systems," in *GLSVLSI*, 2014.
- [43] N. Papandreou *et al.*, "Enhancing the Reliability of MLC NAND Flash Memory Systems by Read Channel Optimization," *ACM Transactions on Design Automation of Electronic Systems*, 2015.
- [44] K.-T. Park *et al.*, "A Zeroing Cell-to-Cell Interference Page Architecture with Temporary LSB Storing and Parallel MSB Program Scheme for MLC NAND Flash Memories," *IEEE Journal of Solid-State Circuits*, 2008.
- [45] T. Parnell *et al.*, "Modelling of the Threshold Voltage Distributions of Sub-20nm NAND Flash Memory," in *GLOBECOM*, 2014.
- [46] A. Prodromakis *et al.*, "MLC NAND Flash Memory: Aging Effect and Chip/Channel Emulation," *Microprocessors and Microsystems*, 2015.
- [47] B. Schroeder *et al.*, "Flash Reliability in Production: The Expected and the Unexpected," in *FAST*, 2016.
- [48] K.-D. Suh *et al.*, "A 3.3 V 32 Mb NAND Flash Memory with Incremental Step Pulse Programming Scheme," *IEEE Journal of Solid-State Circuits*, 1995.
- [49] Terasic, Inc., "Altera DE0 Board," 2013. [Online]. Available: <http://de0.terasic.com/>
- [50] G. Tressler *et al.*, "Enterprise MLC NAND Industry Comparison," in *Flash Memory Summit*, 2011.

- [51] J. P. van Zandwijk, "A Mathematical Approach to NAND Flash-Memory Descrambling and Decoding," *Digital Investigation*, 2015.
- [52] H. Yang *et al.*, "Reliability Issues and Models of Sub-90nm NAND Flash Memory Cells," in *ICSICT*, 2006.
- [53] J. H. Yoon and G. A. Tressler, "Advanced Flash Technology: Status, Scaling Trends, and Implications to Enterprise SSD Technology Enablement," in *Flash Memory Summit*, 2012.
- [54] C. Zambelli *et al.*, "Reliability Issues of NAND Flash Memories," in *Inside NAND Flash Memories*. Springer Netherlands, 2010.
- [55] L. Zhang *et al.*, "Identification of NAND Flash ECC Algorithms in Mobile Devices," *Digital Investigation*, 2012.
- [56] K. Zhao *et al.*, "LDPC-in-SSD: Making Advanced Error Correction Codes Work Effectively in Solid State Drives," in *FAST*, 2013.