# How I Learned to Stop Worrying and Love Flash Endurance

Vidyabhushan Mohan, Taniya Siddiqua, Sudhanva Gurumurthi, Mircea R. Stan†
*Dept. of Computer Science, †Dept. of Electrical and Computer Engg., University of Virginia*
{*vm9u,taniya,gurumurthi,mircea*}*@virginia.edu*

## Abstract

Flash memory in Solid-State Disks (SSDs) has gained tremendous popularity in recent years. The performance and power benefits of SSDs are especially attractive for use in data centers, whose workloads are I/O intensive. However, the apparent limited write-endurance of flash memory has posed an impediment to the wide deployment of SSDs in data centers. Prior architecture and system level studies of flash memory have used simplistic endurance estimates derived from datasheets to highlight these concerns. In this paper, we model the physical processes that affect endurance, which include both stresses to the memory cells as well as a recovery process. Using this model, we show that the recovery process, which the prior studies did not consider, significantly boosts flash endurance. Using a set of real enterprise workloads, we show that this recovery process allows for orders of magnitude higher number of writes and erases than those given in datasheets. Our results indicate that SSDs that use standard wear-leveling techniques are much more resilient under realistic operating conditions than previously assumed and serve to explain some trends observed in recent flash measurement studies.

## 1 Introduction

Flash memory has gained tremendous popularity in recent years. Although initially used only in mobile devices, the drop in the price of NAND flash memory has paved the way for its use in mass storage devices as well, in the form of Solid State Disks (SSDs). SSDs offer several advantages over Hard Disk Drives (HDDs) such as lower power, higher I/O performance (especially for random I/O), and greater ruggedness.

Despite these benefits, one of the main impediments to the wide adoption of SSDs in servers has been its apparent limited write endurance. Flash memory blocks can wear out after a certain number of write (program) and erase operations. Manufacturer datasheets quote values that range from 10,000-100,000 program/erase (P/E) cycles for NAND flash endurance. Architecture and systems papers that explore the use of flash memory use these values to estimate endurance [1]. However, at the physical level, endurance is a more complex process, involving *stresses* due to charge trapping in the tunnel oxide of the floating gate transistors induced by P/E operations, but also a *recovery* process that detraps the charge and partially heals the devices [18, 11]. By modeling both stress and recovery, we can get deeper insights into the endurance characteristics of flash and make a more accurate assessment of SSD endurance in enterprise storage systems. Recent papers on NAND flash chip measurements provide evidence that endurance is higher than the values reported in datasheets [6, 2], which further motivates studying this phenomenon in more detail.

In this paper, we make the following contributions:

- We develop an analytical endurance model for NAND flash memory suitable for use in architecture and system design research.
- We use this model to quantify the impact of charge trapping and detrapping for both single-level cell (SLC) and multi-level cell (MLC) NAND flash.
- We study the endurance of an enterprise-class SSD when exercised by real enterprise workloads. We show that, due to charge detrapping, NAND flash that uses standard wear-leveling techniques can support *two orders of magnitude higher* number of P/E cycles than those given in datasheets. These results indicate that SSDs can be safely deployed in data centers without endurance concerns and explains the reason behind some surprising observations in recent flash measurement studies [6, 2].

The organization of the rest of the paper is as follows. The next section explains flash memory operation and how these operations affect endurance. We then present our analytical endurance model in Section 3 and the impact of charge detrapping is quantified in Section 4. Section 5 describes the experimental methodology for the SSD-level endurance experiments and the results from these experiments are given in Section 6. Finally, Section 7 explains the future work and concludes the paper.

## 2 Flash Memory Operation and Flash Reliability

This section provides an overview of how NAND flash memory operates and explains how these operations affect flash endurance. A detailed discussion on flash memory at the circuit level is given in [3] and [1] describes the architecture of flash based SSDs. Flash is a type of EEPROM (Electrically Erasable Programmable Read-Only Memory) which supports three basic operations: read, program (write), and erase. A flash memory chip consists of the flash memory array and additional peripheral circuitry to perform operations. The flash memory array consists of Floating Gate Transistors (FGTs), which act as memory cells (in this paper, the terms "memory cell" and "FGT" refer to the same physical entity and are used interchangeably). The FGT is similar to a regular MOS transistor except for an additional floating gate between the channel and the control gate. This floating gate is isolated from the rest of the device by dielectric (oxide). This helps retain charges on the floating gate for an extended period of time (on the order of years), hence providing non-volatility. Adding or removing charges to/from the floating gate causes a shift in the threshold voltage of the FGT and this shift is sensed during read. The memory array is partitioned into *blocks* that are, in turn, subdivided into *pages*. A page is the smallest granularity at which the read and program

operations are performed. NAND flash does not support in-place writes and hence an erase operation is necessary before reprogramming the page. Such space management tasks within a SSD are performed by a Flash Translation Layer(FTL).

These program and erase operations are *stress events* that have a detrimental impact on the reliability of flash memory as they affect both *retention* and *endurance*. The typical data retention time for flash memory is 10-20 years [9]. However, as a flash memory cell is repeatedly programmed and erased, the oxide layer becomes weak which leads to an increase in the Stress Induced Leakage Current (SILC) of the memory cell, thus affecting data retention. On the other hand, endurance is a measure of the number of P/E cycles that a flash memory cell can tolerate while preserving the integrity of the stored data, and is a function of the charge trapping characteristics of the oxide [18, 11]. Every stress event increases the likelihood of charges getting trapped in the oxide, which can lead to an undesirable increase in the threshold voltage of the memory cell. If a sufficiently high number of charges get trapped in the oxide, it will no longer be possible to reliably read the cell.

Although a memory cell that undergoes a large number of stress events will have more charges trapped in its oxide, several transistor-level studies of NAND-flash memory have shown that it is possible to *detrap (i.e., remove)* some of the charges from the tunnel oxide under certain conditions [18, 11, 17]. Beneficial conditions for detrapping include higher external temperatures and quiescent periods between successive stress events. Furthermore, the measurement studies indicate that introducing a quiescent period to allow detrapping can be applied at temperatures as low as $25°C$, which is the typical external ambient temperature of a disk [7]. Since the quiescent periods help improve endurance, we refer to them as *recovery periods*.

# 3 Flash Endurance Model

In order to analyze how stress events and recovery periods impact the endurance of NAND flash memory under various usage scenarios, we have developed an analytical model that captures how these two parameters affect the threshold voltage of memory cells. This model is constructed by synthesizing the results from device physics papers on NAND flash memory cells [18, 11, 19]. These papers provide information about how the parameters are related and also provide memory cell level measurements of stresses and recovery for a range of values.

The model consists of two parts - one for stresses and the other for recovery. The first part of the model gives the relationship between the increase in threshold voltage due to charge trapping ($\delta V_{th,s}$) and the number of stress events on the oxide. The second part gives the relationship between the threshold voltage shift due to recovery ($\delta V_{th,r}$), the amount of trapped charges in the oxide due to stress calculated in the first part ($\delta V_{th,s}$), and the recovery period (t). Using these two parts, we calculate the effective increase in threshold voltage of a memory cell due to trapped charges ($\delta V_{th}$) after a stress event and a subsequent recovery period. We now explain these two components of the model in more detail.

## 3.1 The Stress Model

The threshold voltage of a memory cell increases due to charge trapping with the number of stress events (program or erase cycles) [19]. There are two types of traps that form in the oxide - interface traps and bulk traps - which contribute to the increase in the threshold voltage. It has been shown that both types of traps have a power-law relation to the number of P/E cycles on the memory cell as [19]:

$$\begin{aligned} \delta N_{it} &= A * cycle^{0.62} \\ \delta N_{ot} &= B * cycle^{0.30} \end{aligned}$$

where $A$ and $B$ are constants, $cycle$ is the number of program or erase cycles on the cell, and the terms $\delta N_{it}$ and $\delta N_{ot}$ are the interface and bulk trap densities respectively. In addition to providing this power-law relationship, [19] also provides empirical data on how $\delta N_{it}$ and $\delta N_{ot}$ vary with $cycle$. We calculated the values of constants $A$ and $B$ to be 0.08 and 5 respectively for the sub-90nm process technology from this empirical data. Similar device characterization data can be used for other process technologies in our model to estimate endurance.

The total threshold voltage increase due to trapping is divided into interface trap voltage shift ($\delta V_{it}$) and bulk trap voltage shift ($\delta V_{ot}$). Park et al. [14] give the relationship between $\delta V_{it}$ and $\delta N_{it}$ and between $\delta V_{ot}$ and $\delta N_{ot}$ to be:

$$\delta V_{it} = \frac{\delta N_{it} * q}{C_{ox}} \quad (1)$$

$$\delta V_{ot} = \frac{\delta N_{ot} * q}{C_{ox}} \quad (2)$$

where $q$ is electron charge ($1.6 \times 10^{-19}$ Coulombs) and $C_{ox}$ is the capacitance of the oxide. The value of $C_{ox}$ depends on the feature size of the NAND flash cell.

Hence the increase in threshold voltage of the memory cell due to trapped charges, $\delta V_{th,s}$, is given by:

$$\delta V_{th,s} = \delta V_{it} + \delta V_{ot} \quad (3)$$

## 3.2 The Recovery Model

According to Yamada et al. [18], the threshold voltage shift due to detrapping depends on the recovery period and the amount of charge trapped in the oxide. This relationship is given by:

$$\delta V_{th,r} = c_{vt} * ln(\frac{t}{t_0}) \quad (4)$$

where $t$ is the recovery period between successive stress events to the *same cell* (in seconds), $t_0$ is $1sec$ and $c_{vt}$ depends on the amount of trapped charge ($Q$) present in the oxide. The value of the recovery period, $t$, is assumed to be finite and greater than one second. We conservatively assume that no charge detrapping occurs for recovery periods less than one second. Yamada et al. also show that $c_{vt}$ has a logarithmic dependence on $Q$ [18]. Since $Q$ is directly proportional to the stress voltage, $\delta V_{th,s}$, $c_{vt}$ also has a logarithmic dependence on $\delta V_{th,s}$. Yamada et al. provides empirical plots of how $c_{vt}$ varies with $\delta V_{th,s}$ [18] . Using these plots, we get:

$$c_{vt} = ln(\frac{\delta V_{th,s}}{V_0}) \quad (5)$$

The unit of $c_{vt}$ is in $mV$ and $V_0$ is $1mV$.

Combining equations (4) and (5), the change in the threshold voltage shift due to recovery, $\delta V_{th,pr}$ is given by:

$$\delta V_{th,pr} = ln(\frac{\delta V_{th,s}}{V_0}) * ln(\frac{t}{t_0}) \quad (6)$$

where $\delta V_{th,s}$ is given by equation (3). Since the physical process of recovery is not perfect, we introduce a term $K$ which denotes the efficiency of recovery process. Discussions with the industry [12] indicate that $K$ does not exceed 60%. Hence, equation (6) is modified as:

$$\delta V_{th,r} = \begin{cases} \delta V_{th,pr} & \text{if } \delta V_{th,pr} < \text{K} * \delta V_{th,s} \\ K * \delta V_{th,s} & \text{otherwise} \end{cases} \quad (7)$$

The effective increase in the threshold voltage due to trapped charges after stress and recovery, $\delta V_{th}$, is given by:

$$\delta V_{th} = \delta V_{th,s} - \delta V_{th,r} \quad (8)$$

Equations (3) and (7) can be used to estimate the endurance of a NAND flash memory cell based on the number of stress events (P/E cycles) and the recovery periods that the cell experiences. The stress events and recovery periods for a workload can be tracked using an storage system simulator such as Disksim [5]. While the model captures the impact of stress and recovery on a single memory cell, we track stress events due to program and erase operations at the granularity of a page and block respectively. Also one can use the methodology given above for both SLC and MLC by varying the maximum allowed threshold voltage shift.

## 3.3 Limitations of the Model

Currently, the model has two limitations:

- The model does not capture the impact of temperature on stress and recovery. The constants in our model are estimated from published datapoints for $25°C$, which is approximately the external ambient temperature of a disk drive in a server with a well-designed cooling system [7].
- If the memory cell is used in MLC mode, programming each n-bit value requires a different amount of time [3] and hence the duration of the stress events would be different for different bit values. Moreover, the amount of charge that is trapped in the oxide also depends on the bits to be stored in the cell. Currently, the model does not account for these variations and estimates the impact of stress and recovery in a way that is agnostic to the actual bits stored in the cells.

Despite these limitations, the model captures the primary effects of charge trapping and detrapping on the endurance of NAND flash and is suitable for use in storage system simulations. This model can be used by system architects to estimate how a particular workload would affect the endurance characteristics of an SSD prior to deployment and configure their storage system accordingly.

# 4 Impact of Charge Detrapping on SSD Endurance and Model Validation

Having derived a model for the threshold voltage shift due to stress and recovery, we now analyze the impact of charge detrapping on flash memory cells over different timescales. The goal of this analysis is to ascertain the extent to which charge detrapping can improve the reliability of flash memory cells by delaying endurance related failure and understand how the duration of the quiescent period affects the extent of the recovery.

Before we begin the analysis, we first need to precisely define what "failure" means with respect to endurance. The data stored in a flash memory cell is identified by a specific voltage level. An n-bit MLC has $2^n$ *distinct* voltage levels, each of which corresponds to an n-bit value (an SLC flash cell is merely the case where n=1, which corresponds to two voltage levels - one for a digital "0" and the other for a "1"). Let $\Delta V_{th,spread}$ be the threshold voltage range for a single voltage level in a memory cell and $\Delta V_{th}$ be the difference in voltage between adjacent levels. This is illustrated in Figure 1 for $n = 2$ (2-bit MLC). When the charges trapped in the oxide result in a threshold voltage increase of $\Delta V_{th}$ or higher, it will no longer be possible to clearly distinguish between different voltage levels. As a consequence, it will not be possible to reliably read from or write to the memory cell. We define this situation as a *failure*. We define the *endurance limit* as the total number of P/E cycles before this condition occurs. Once the endurance limit is reached, a page is considered to have failed and is no longer usable. Manufacturer datasheets specify an endurance limit of 10K and 100K P/E cycles for MLC and SLC chips respectively. However, these values specify the *minimum* number of P/E cycles that the chip is expected to tolerate before failure, tested under high stress conditions where the flash cells are continuously erased and rewritten with little or no recovery time between successive stress events [17]. There is anecdotal evidence in recently published papers on measurements of NAND flash chips [6, 2], that, in the common case, when there are recovery periods between the stress events, the endurance of flash is higher than the values specified in datasheets.

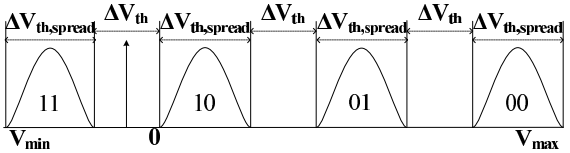


Figure 1: Threshold voltage distribution for a 2-bit MLC

In Figure 2, we plot the change in $\delta V_{th}$ with the number of P/E cycles, over a number of timescales for the recovery period, for the sub-90nm process technology. We consider the case where there is no recovery between successive stress events, which is how the datasheet values are computed, and also cases where the recovery time is varied from 10 seconds to over 2 days. To illustrate how these curves translate to endurance, we plot the $\Delta V_{th}$ for SLC and 2-bit MLC flash. These values are shown as horizontal lines in the graph and are obtained from threshold voltage distributions of prototype NAND flash memory chips published in the literature. 2-bit MLC devices have $\Delta V_{th}$ values that are approximately equal to $\Delta V_{th,spread}$ and have been shown to vary from 0.6V to 0.7V [4]. We assume $\Delta V_{th}$ to be equal to $\Delta V_{th,spread}$ for SLC devices as well. The $\Delta V_{th,spread}$ of SLC has been reported to vary from 1.4V to 2.0V [13, 10]. Based on this data, we assume the $\Delta V_{th}$ of SLC to be 1.7V and 2-bit MLC to be 0.65V. The por-

tions of the curves below the horizontal lines correspond to failure-free operation of the cell. The number of P/E cycles attainable for each recovery period and the improvement in endurance over the case where there is no detrapping between successive stresses is given in Table 1 (for clarity, in the figure we omit a few of the datapoints given in the table).

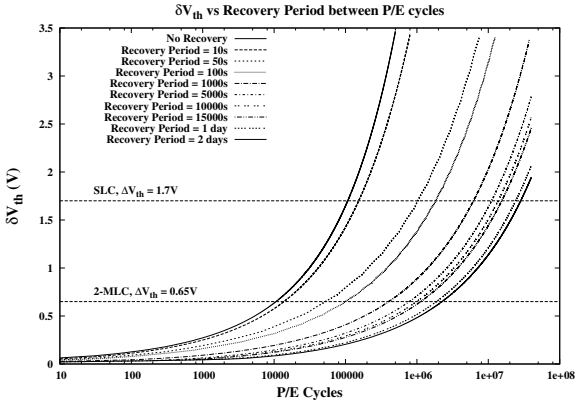


Figure 2: Increase in $\delta V_{th}$ with P/E cycles for different recovery periods.

| Recovery Period | SLC, $\Delta V_{th} = 1.7V$ | | 2-bit MLC, $\Delta V_{th} = 0.65V$ | |
|---|---|---|---|---|
| | P/E Cycles | Endurance Increase | P/E Cycles | Endurance Increase |
| No recovery | 107535 | 1x | 10652 | 1x |
| 10 seconds | 153186 | 1.4x | 13749 | 1.3x |
| 50 seconds | 1028724 | 9.6x | 52444 | 4.9x |
| 100 seconds | 1837530 | 17.7x | 99913 | 9.3x |
| 1000 seconds | 6214983 | 57.8x | 403082 | 37.8x |
| 5000 seconds | 11093823 | 103x | 780723 | 73.3x |
| 10000 seconds | 13753999 | 127x | 990014 | 92.9x |
| 15000 seconds | 15497892 | 144x | 1129379 | 106x |
| 1 day | 24274492 | 225x | 1879352 | 176x |
| 2 days | 28487539 | 264x | 2247910 | 211x |

Table 1: Endurance limits with charge detrapping.

*We can see that, when there are no recovery periods, the P/E cycles for the SLC and MLC datapoints approximately match the values given in datasheets (100K and 10K P/E cycles respectively), which concurs with the expected behavior.* We can also see that a recovery period between successive P/E cycles can significantly boost endurance, which concurs with recent flash chip measurement studies [6, 2]. A recovery period of a few hours provides two orders of magnitude improvement. However, as the recovery periods increase beyond a day, we start getting diminishing endurance benefits.

**The Role of ECC:** It is important to note that charge detrapping does not preclude the use of ECC. ECC is still required to handle bit errors that happen due to disturb events and hence complements detrapping to boost reliability.

## 5 Experimental Methodology

**Simulator and Workloads:** We use the Disksim SSD extension [1] that facilitates studying a variety of SSD designs. We simulate a 32GB SSD composed of 8 4GB SLC NAND flash chips, similar to enterprise-class SSDs currently available [8].

Our workloads consist of block-level I/O traces collected from various production systems within Microsoft [16, 15]. We evaluate four workloads: Live Maps™ (LM), Exchange™(EXCH), RADIUS™ (RAD), and MSN™ (MSNFS). Each workload consists of several sub-traces, each of which correspond to the I/O activity during a specific interval of time (e.g., an hour) on a typical day, and the collection of these sub-traces span at least one full day. We use all the sub-traces of each workload in the simulation to characterize the variations in the I/O behavior and their impact on the SSD over a one-day period.

**Endurance Metric:** We report endurance in terms of the number of P/E cycles. Since different blocks may undergo a different number of P/E cycles based on the workload and FTL behavior, we report the average number of P/E cycles and the minimum and maximum values observed across all the blocks. We report endurance at the end of 5 years of activity. This 5-year service life allows us to examine the impact of P/E cycles on endurance well before retention related reliability problems arise (the retention period of flash is 10-20 years [9]).

**Estimating Endurance Over the Service Life:** Since the service life spans multiple years whereas the traces record only a single day of activity, we need a way of estimating the activity on the SSD over this long time period. Since each trace represents the I/O activity over the course of a typical day, one approach could be to repeatedly replay the trace in Disksim and simulate 5 years worth of activity. However, this approach would require excessively long simulation times. We instead use a statistical approach to estimate the I/O activity on the SSD.

In order to estimate endurance, we need to capture two aspects of stress behavior: (1) the distribution of stress events across various pages and blocks in the SSD (spatial behavior), and (2) the distribution of the recovery periods to individual pages and blocks (temporal behavior). To determine these distributions, we collect an output trace over the course of a Disksim simulation that records when a particular page or block within a certain flash chip is programmed or erased. We collect one such output trace for each sub-trace,which allows us to capture any phase behavior within a workload. From this output trace, we characterize the spatial behavior of the workload by creating a histogram of the stresses to the different flash chips in the SSD to determine the frequency at which pages/blocks within a particular chip are stressed. Since wear-leveling operations are performed within each flash chip in an SSD [1] we use a uniform distribution to model the pattern of stresses within a chip. We characterize the temporal behavior of the workload by creating a histogram of the recovery periods of all the pages within the SSD. Using these statistical distributions of the spatial and temporal characteristics of a workload's stress behavior on the SSD, we extrapolate the stress behavior over the service life of 5 years.

## 6 Results

The number of P/E cycles over the 5-year service life of the SSD for each workload is given in Figure 3(a). Each bar gives the average number of P/E cycles across all the SSD blocks while the error-bar shows the smallest and largest number of P/E cycles observed across the blocks. We can see that there is significant variation in terms of the number of P/E cycles that blocks experience across the different workloads over the SSD service life. The RAD and EXCH workloads impose fewer stresses on the blocks whereas LM and MSNFS impose far greater number of stresses. The change in $\delta V_{th}$ for the block that experienced the largest number of P/E cycles for

each workload is given in Figure 3(b). We observe that across all the workloads, the increase in $\delta V_{th}$ over the 5-year period is well below the SLC $\Delta V_{th}$ of 1.7V. This is true even for LM and MSNFS whose P/E cycles are close to or greater than the datasheet specified endurance limit of 100K P/E cycles. This is due to detrapping during the recovery periods between stress events. We can also observe that, although the number of P/E cycles for the most heavily stressed block in EXCH is significantly lower than that in MSNFS, the $\delta V_{th}$ of MSNFS is only slightly higher than that of EXCH. This is because the relationship between $\delta V_{th}$ and the number of P/E cycles is not linear, as discussed in Section 3.

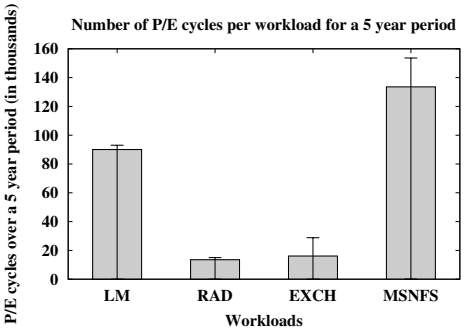


(a) P/E cycles experienced over the service life. The error-bars correspond to the smallest and largest number of cycles for a block in the SSD.

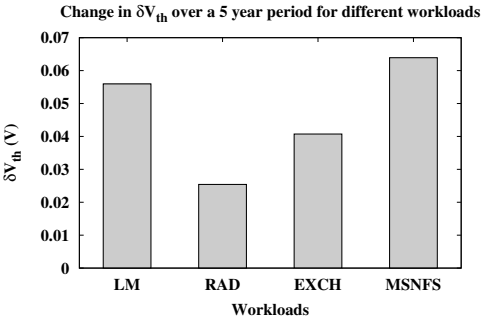


(b) Change in $\delta V_{th}$ for the block with the highest P/E cycles. $\Delta V_{th}$=1.7V for SLC Flash.

Figure 3: Endurance results for enterprise workloads.

The distribution of the recovery periods of the SSD blocks for the workloads over the service life is given in Table 2. We can see that most, if not all, SSD blocks in all four workloads experience recovery periods in the order of thousands of seconds. As Table 1 indicates, recovery periods of such durations can significantly boost endurance, allowing the blocks to undergo several *millions* of P/E cycles before reaching the endurance limit. The amount of time required for reaching the endurance limit is much longer than the NAND flash retention period. Therefore, endurance is not a major flash reliability concern under realistic data center usage scenarios and a much wider array of I/O intensive applications can leverage the performance and power benefits of flash-based SSDs than previously assumed.

| Benchmarks | Recovery Period (seconds) | | | |
|---|---|---|---|---|
| | [1k-5k) | [5k-10k) | [10k-15k) | [15k-20k) |
| **LM** | 100 | 0 | 0 | 0 |
| **RAD** | 0.001 | 0.0023 | 99.9957 | 0.001 |
| **EXCH** | 0.002 | 99.857 | 0.141 | 0 |
| **MSNFS** | 100 | 0 | 0 | 0 |

Table 2: Recovery time distributions. $[X, y)$ indicates recovery periods of duration $t$ where $X \leq t < y$.

# 7 Future Work

While this work has focused on one aspect of flash memory reliability, namely endurance, we are also trying to answer other questions related to flash memory reliability. Our immediate focus is to validate our model with real chip measurements. We are in the process of building a test board for this validation. As mentioned in Section 3.3, our model does not capture the impact of temperature on stress and recovery. We plan to model the impact of this parameter in the future. Another aspect of reliability we plan to model is Bit Error Rate (BER). Prior studies have shown that BERs depend on the age of a flash chip [6]. Modeling the relation between the age of a flash chip and BERs will provide insights into the strength of ECC required for correcting such errors. We also plan to model SILC to factor-in retention. Overall, modeling these phenomena provides a comprehensive analysis of NAND flash reliability.

# 8 Acknowledgments

We thank Neal R. Mielke (Intel), Sriram Sankar (Microsoft), Steve Schlosser (our shepherd), and the anonymous reviewers for their valuable inputs. This research has been supported in part by NSF CAREER Award CCF-0643925, NSF grant CNS-0551630, and gifts from Google and HP.

# References

[1] N. Agrawal and et al. Design Tradeoffs for SSD Performance. *USENIX*, 2008.
[2] S. Boboila and P. Desnoyers. Write Endurance in Flash Drives: Measurements and Analysis. *FAST*, 2010.
[3] J. Brewer and M. Gill, editors. *Nonvolatile Memory Technologies with Emphasis on Flash*. IEEE Press, 2008.
[4] T. Cho and et al. A dual-mode NAND flash memory: 1-Gb multilevel and high-performance 512-Mb single-level modes. *JSSC*, 2001.
[5] G. Ganger, B. Worthington, and Y. Patt. *The DiskSim Simulation Environment Version 4.0 Reference Manual*.
[6] L. Grupp and et al. Characterizing flash memory: Anomalies, observations, and applications. *MICRO 2009*.
[7] S. Gurumurthi, A. Sivasubramaniam, and V. Natarajan. Disk Drive Roadmap from the Thermal Perspective: A Case for Dynamic Thermal Management. *ISCA*, 2005.
[8] Intel X25-E Extreme SATA Solid-State Drive.
[9] Process Integration and Device Structures, ITRS 2007 Edition.
[10] J. Lee and et al. A 1.8V NAND Flash Memory for Mass Storage Applications. *ISSCC*, 2003.
[11] N. Mielke and et al. Recovery effects in the distributed cycling of flash memories. *IRPS*, 2006.
[12] N. R. Mielke. Intel Corporation, October 2009. Private Correspondence.
[13] H. Nakamura and et al. A 125mm$^2$ 1Gb NAND Flash Memory with 10 MB/s Program Throughput. *ISSCC*, 2002.
[14] Y. Park and D. Schroder. Degradation of thin tunnel gate oxide under constant fowlernordheim current stress for a flash eeprom. *TED*, 1998.
[15] S. Kavalanekar and et al. Characterization of storage workload traces from production Windows servers. *IISWC*, 2008.
[16] SNIA IOTTA Trace Repository.
[17] Application Report: MSP430 Flash Memory Characteristics.
[18] R. Yamada and et al. A novel analysis method of threshold voltage shift due to detrap in a multi-level flash memory. *VLSI Symposium*, 2001.
[19] H. Yang and et al. Reliability issues and models of sub-90nm nand flash memory cells. *ICSICT*, 2006.