

Underpowering NAND Flash: Profits and Perils

Hung-Wei Tseng Laura M. Grupp Steven Swanson
The Department of Computer Science and Engineering
University of California, San Diego
{h1tseng,lgrupp,swanson}@cs.ucsd.edu

Abstract

MLC Flash memory is getting more popular in computer systems ranging from sensor networks and embedded systems to large-scale server systems. However, MLC flash has many reliability concerns, including the potential for corruption due to supply voltage fluctuations. This paper characterizes MLC flash when the chip is underpowered (i.e., power fading and voltage droops). We demonstrate that underpowering flash can cause serious errors, but also help saving up to 45% of operation energy without incurring failure.

1. INTRODUCTION

In recent years, computer systems ranging from sensor networks to portable devices to data centers have embraced flash-based solid-state drives (SSDs) to provide low power, light weight and shock resistant storage. These applications can expose flash memories to a range of stresses including unstable or unreliable power supplies that can underpower flash memories.

Flash-based SSDs have several characteristics that make unstable or unreliable power supplies particularly dangerous. For instance, an inopportune droop in supply voltage could corrupt flash translation layer (FTL) metadata and render the whole storage device unusable. If designers hope to implement reliable storage systems, they must understand the behavior of flash devices under these conditions.

This paper characterizes the behavior of flash memory devices when supply voltage droops unexpectedly during an operation. We show that supply voltage droops can incur errors for all operations. We also demonstrate that, for some operations, significantly reducing the supply voltage does not impact reliability. In these cases, we can exploit the chips' tolerance for drooping voltages to save energy.

To characterize the behavior of flash memory chips, we developed a testing platform to repeatedly change the supply voltage to a raw flash device during read, program and erase operations. We examined 5 popular MLC chips from different vendors. Our data show several interesting behaviors of underpowering flash memory chips. First, underpowering flash memory during an operation does cause data corruption in some cases, but not always lead to incomplete operations. Second, the minimum voltage each operation requires to complete an operation without increasing the error rate is different. Third, underpowering flash chips can negatively impact the latency of operations but positively improve the energy consumption. We utilize our characterization results to design a dy-

namic voltage scaling mechanism that adjusts the supply voltages to flash chips according to the type of operations. We can achieve up to 45% energy saving for the flash storage device without the assistance of any special data encoding scheme.

The rest of this paper is organized as follows: Section 2 describes the aspects of flash memory pertinent to this study. Section 3 describes our experimental platform and methodology for characterizing flash memory's behavior when supply voltage droops. Section 4 presents our characterization results. Section 5 discusses the potential energy saving using our characterization result. Section 6 provides a summary of related work to put this project in context, and Section 7 concludes the paper.

2. FLASH MEMORY

This section presents a brief introduction of flash's characteristics that are most relevant to understanding its behavior in the face of unstable power supplies.

A cell in the flash array stores data by trapping electrons using a floating gate transistor. The electrical charge of the floating gate affects the transistor's voltage level, and the chip compares this level voltage and with threshold voltages to read the data that the cell currently stores.

According to the number of voltage levels of each cell, flash cells can divide into single-level cell (SLC) and multi-level cell (MLC). SLC devices store one bit per cell, while MLC devices store two or more. MLC chips obtain higher densities by representing n bits using 2^n threshold voltage levels. SLC chips provides better and more consistent performance and lifetime than MLC chips. The empirical measurements in [3] show that MLC chips need $30 \mu\text{s}$ – $50 \mu\text{s}$ to perform a read operation, $300 \mu\text{s}$ – 2ms to perform a program operation, and 2ms – 4ms to perform an erase operation. In this paper, we focus on 2-bit MLC cells, since they are most prevalent in current systems.

The flash chip divides cell arrays into blocks (between 32 and 256 pages) and blocks into pages (between 2 and 8 KB). The erase operation sets all the bits in a block to '1'. The erase operation operates at the granularity of one block. The program operation converts 1s to 0s and performs operates at page granularity. Each block in flash memory has limited number of erase cycles (lifetime). Due to the limited lifetime and the difference in granularity between programs and erases, SSDs apply complex flash translation layers (FTLs) to perform out-of-place update and remapping operations to improve performance and lifetime. FTLs need to store metadata in the flash storage array along with the user data.

The erase operations is iterative. The chip removes the electrical charges from cells within the block and checks if all the cells reaches the erased state. The chip will continue to remove electrons until the threshold voltage levels of cells reach the erased state.

The program operation injects electrons into floating gates to change the threshold voltage for the cells and then perform a read-verify operation to check if the cells have reached the target threshold voltage. If any of the cells in a page failed to pass the check, the chip will repeat the program and the read-verify process [13, 7].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DAC'13, May 29 - June 07, 2013, Austin, TX, USA.

Copyright 2013 ACM ACM 978-1-4503-2071-9/13/05 ...\$15.00.

Voltage Levels	Logic Bits			
	Gray coding		2's complement coding	
	1st page bit	2nd page bit	1st page bit	2nd page bit
Lowest	1	1	1	1
	1	0	1	0
	0	0	0	1
Highest	0	1	0	0

Table 1: The mapping of voltage level and logic bits in 2-bit MLC chips using Gray coding and 2's complement coding.

For 2-bit MLC chips, each cell stores one bit from two different, logical pages for two different pages. Manufactures require that pages within a block be programmed in order, so to differentiate between the two pages in a cell, we refer to them as “first page” and “second page”. Programming a first page only sets the threshold voltages of cells to contain one bit data. Programming a second page further divide the existing 1s and 0s into four different states to present 2-bit data. Table 1 shows the mappings between threshold voltages and logic bits of a 2-bit MLC cell using gray coding and 2's complement coding after the chip programmed a second page of a cell. Programming a second page takes longer time than programming a first page, because programming the second page requires a more complex process.

3. METHODOLOGY

To study behavior flash memory under an unstable power supply, we designed and implemented a test system that can issue commands to raw flash chips and change the supply voltage to the chip precisely. This section presents our test system, experimental methodology, and the flash chips we examined in this study.

3.1 Experimental hardware

Our test system contains a Xilinx XUP board, a custom flash testing board, and the voltage control circuit. The Xilinx XUP FPGA board implements a custom flash controller. The PowerPC core on the FPGA board allows the system to host a full-fledged Linux distribution. The system also provides libraries that allow benchmark applications to directly access the flash device on the flash testing board through our custom controller and control the supply voltage. The flash testing board provides sockets for installing flash devices and facilities for measuring power and performance. The flash testing board uses 3.3 V as the default supply voltage. The voltage control circuit consists of high-speed transistors and accepts signals from the FPGA board to change the supply voltage to the flash chips. We use a high-resolution oscilloscope to measure the energy consumption of operations. To minimize switching time, we remove all capacitors on the flash testing board. The oscilloscope shows that the resulting system can change the chip's supply voltage within 10 ns and the time takes to reach a stable voltage level is slightly less than 2 μ s. We also examined the behavior of flash memory devices between 2.7 V – 3.3 V with the capacitors. Removing the capacitors does not result in significant differences in bit error rates and latencies.

3.2 Test procedure

To examine the behavior of underpowering flash memory chips during read, program and erase operations, we change supply voltage of flash memory chips from the suggested operating voltage (2.7 V – 3.3 V) to a voltage below the chip's specified minimum operating voltage (1.8 V – 2.6 V) during the operations. We define the *voltage change interval* as the time between issuing a command to a flash device and when the FPGA board triggers the change of supply voltage. The voltage change interval starts after the chip receives the last byte of the command. The high-resolution oscilloscope shows that the chip starts executing the command within 10 μ s.

Abbrev.	Manufacturer	Cell Type	Cap. (Gbit)	Tech. Node (nm)	Page Size (B)	Pgs/Blk
A-MLC16	A	MLC	16		4096	128
B-MLC32	B	MLC	32	34	4096	256
B-MLC128	B	MLC	128	34	4096	256
E-MLC8	E	MLC	8		4096	128
F-MLC16	F	MLC	16	41	4096	128

Table 2: Parameters for flash devices we studied in this work

For read operations, we use voltage change intervals varying from 0.4 μ s to 60 μ s at increments of 0.4 μ s. For program tests, we use voltage change intervals varying from 0.4 μ s to 2.4 ms at increments of 0.4 μ s. For erase, we use voltage change intervals varying from 2 μ s to 4.8 ms at increments of 2 μ s.

3.3 Flash devices

Flash memory chips from different manufacturers demonstrate various behaviors because of their architectural differences within the devices and manufacturing technologies. In this work, we selected 5 MLC chips that cover a variety of technologies and capacities to better understand the variation of flash devices when voltage droop occurs. According to the data-sheets, all these chips have suggested operation voltage range between 2.7 V and 3.6 V. Each chip that we used in this project guarantees 5000 erase cycles for their blocks. We perform all our experiments under the suggested lifetime of blocks.

Table 2 illustrates the flash memory chips that we studied in this work. These devices come from five different manufacturers with process technologies ranging from 72 nm to 34 nm. Their capacities range from 8 Gbits to 128 Gbits. We obtain values that are not publicly available from the manufacturer from [3].

4. EXPERIMENTAL RESULTS

We found interesting behavior when the supply voltage droops during operations. It appears that the flash memory chips can tolerate the changing of supply voltage within a certain range and allow operations to complete without incurring additional errors. The voltage ranges that each chip can tolerate are different for different type of operations. The supply voltage also affects the latencies and energy of flash operations. We describe the details of our results in this section.

4.1 Program and power fade

To understand the impact of underpowering flash chips during program operations, we begin by programming each page in a block with random data and changing the supply voltage from the 3.3 V to a lower voltage level (1.8 V – 2.6 V) at different voltage change intervals. Then, we measure the resulting bit error rate and latency.

We found that all chip behave similarly when the supply voltage droops during program operations. First, lowering the supply voltage during a program operation does not always incur more errors. Second, lowering the supply voltage will increase the latency to complete a program operation. Finally, we find that lowering the voltage during the program operation can potentially corrupt data already present in flash memory.

4.1.1 Programming pages

To illustrate our experimental result, we use E-MLC8 as an example. Figure 1 shows the average bit error rates when the supply voltage droops from 3.3 V to 1.8 V – 2.6 V at different voltage change intervals. When the target supply voltage is above 2.1 V, the bit error rates do not show significant difference comparing with performing the whole operation at 3.3 V. When the voltage droops to 2.0 V, the bit error rate curve shows some spikes indicating that




Figure 1: The average bit error rate for E-MLC8 when the supply voltage droops to 1.8 V – 2.6 V at different voltage change intervals.




Figure 2: The average program latency for E-MLC8 when the supply voltages droops to 1.8 V – 2.6 V at different voltage change intervals

the bit error rate can suddenly increase at a certain voltage change intervals because the insufficient power supply corrupts the process of changing threshold voltages in some cells. Underpowering the chip to 2.0 V during programming is dangerous. If we change the supply voltage to 1.9 V or 1.8 V within 4 μ s of the operation starting, the bit error rate reaches 50%, which means the program operation completely failed. For voltage change interval between 4 μ s and 480 μ s, some cells fail to finish programming their first page bits and all cells fail to program the second page bits resulting in an 37.5% bit error rate during this period. At 480 μ s, the chip applies a different threshold voltage to distinguish the second page data, and the bit error rate drops significantly. However, the bit error rates of the second page bits remain high until the voltage change interval reaches 556 μ s.

We also found that lowering the supply voltage will increase the program latency for both first and second pages. Figure 2 shows the average latency of E-MLC8 when the supply voltage droops from 3.3 V to 1.8 V – 2.6 V. The latencies for voltages 1.8 V and 1.9 V shows very low latency when the voltage change interval is smaller than 556 μ s. The chip aborts the program and read-verify process because the chip cannot reprogram and correct these high bit error rates using the lower power supply. When the supply voltage droops to 2.1 V – 2.6 V during the operations, the program latency increases. In general, lower voltages result in longer latencies. For 2.1 V, which is the minimum supply voltage that allows the program operation to complete without incurring significant program error rates, the latency increases by 20% if we change the voltage immediately after the operation begins.

Programming a flash page with lower supply voltage increases the latency but can help reduce the energy consumption. For example, in E-MLC8, programming at 2.1 V can reduce the energy consumption by 32% for a first page, and 34% for a second page.

4.1.2 Retroactive data corruption




Figure 3: The retroactive data corruption problem occurs when we change the supply voltage for E-MLC8 during a program operation

The retroactive data corruption [14] is a phenomenon that a later incomplete program operation can corrupt data that the chip already successfully programmed. The prior work demonstrates this phenomenon can corrupt 50% of first page bits if power failure occurs during second page program operations. We also examined this problem when the supply voltage droops during a program operation.

Figure 3 illustrates this effect for E-MLC8. In this graph, we program random data into the first page bits of E-MLC8 without changing the supply voltage. However, when we program the second page, we change the supply voltage to 1.8 V during the operations at different voltage change intervals. The x-axis shows the voltage change intervals for the second page program operations, and the y-axis shows the bit error rates of the first pages after we programmed the second page.

For E-MLC8, the bit error rate reaches 25% if the voltage droop happens between 36 μ s and 556 μ s during the second page program operation. This result indicates that SSDs can potentially lose data that are already present in the chip if power supply is unstable.

4.1.3 Read disturb sensitivity

Another data integrity issue that we examined is the relationship of read disturb and power fade. Read disturb happens because the read operation of flash array applies weak programming voltages to cells in pages not involved in the read operation. It typically takes several million read operations to cause significant errors due to read disturb [3]. In this work, we examine the read disturb sensitivity if we underpowering flash chips during program operations.

To examine the read disturb problem for programming flash pages with low voltages, we program flash pages within the chip using the lowest voltage that does not incur more errors (e.g. 2.1 V for E-MLC8) than programming at 3.3 V for any voltage change interval.

Our result shows that for programming operation finished with the selected voltage level, the read disturb sensitivity does not show any difference comparing with programming at normal voltage.

4.1.4 Data retention

The previous experiments demonstrate that program operation can potentially complete with lower supply voltage. In this section, we also examine the long-term stability of the programming data under lower supply voltages. We use a laboratory oven to bake the chips for 9.33 hours at 125°C to simulate the aging effect of one year [5], and repeat the same process until each chip reaches the aging of 10 years. For each chip, we programmed 5 blocks for two conditions: (1) Programming at minimum program voltage: We program flash pages using the lowest voltage that does not incur more errors (e.g. 2.1 V for E-MLC8) than programming at 3.3 V for the whole operation. (2) Programming at normal voltage: the

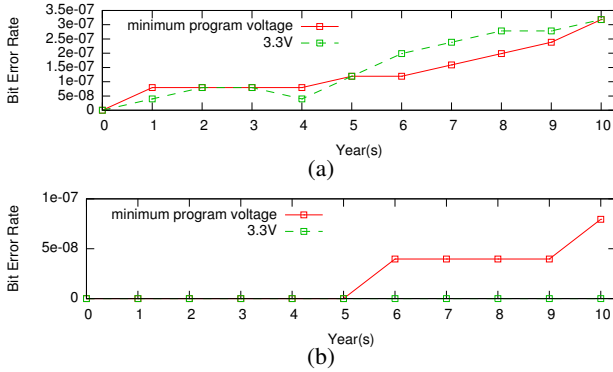


Figure 4: Accelerating aging experiments reveal that program flash chips using lower supply voltage can affect the long-term reliability.

program operation completes using 3.3 V.

The experimental result falls into two categories. We use two chips, F-MLC16 (Figure 4(a)) and A-MLC16 (Figure 4(b)), as examples. The x-axes of Figure 4 are accelerated age in years. The y-axes show the average bit error rates in a block after aging for each accelerated year. For F-MLC16, the data retention of programming at minimum program voltage does not exhibit any difference from programming at normal voltage. For A-MLC16, programming at lowest program voltage results in 1 to 3 bits errors in each page with an average error rate of 7.95×10^{-8} rather than 0 for programming at normal voltage. Though lowering the programming voltage can potentially hurt data retention, the error rates are still manageable with usual error correction codes.

4.2 Read and power fade

To examine the behavior of flash memory chips when the supply voltage droops during read operations, we first program each page in a block with random data at the normal voltage level (3.3 V). Then, we read back the data but change the supply voltage during the operations and measure the bit error rate and the latencies.

The behavior for read operation shows some diversity across chips. For B-MLC32 and B-MLC128, these two chips can never safely read beyond the recommend operating voltages. Take Figure 5(a) as an example, the chip hits 25% error rate when the voltage droops to 1.9 V. If the target voltage is between 2.1 V and 2.6 V, the chip seems to work properly when the voltage change intervals is less than $10 \mu s$. But the read operation, again, shows significant errors if the voltage change intervals are more than $10 \mu s$. Reading these flash memory chips at any voltage below the suggested range is dangerous. There is a cliff at $42 \mu s$ because the read operation completes before the voltage changes.

For the rest of the chips we examined, the behaviors are similar to E-MLC8 in Figure 5(b). For E-MLC8, if the voltage droops to below 2.2 V, the chip will show significant errors no matter when we switch the supply voltage. However, if the voltage only droops to 2.5 V, which is still below the suggested range, the bit error rate is as low as reading this chip at 3.3 V. Comparing with program operations, the read operations for all chips that we examined require higher supply voltages. The minimum voltages that allow read operations to complete without incurring significant errors are also very close to the lower bound of the suggested operation voltage.

Lowering the supply voltage also increases the latency across all chips we tested. Figure 6 shows the read latencies of E-MLC8. For E-MLC8, dropping the supply voltage to 2.5 V during a read operation increases the latency by 1%.

As in program operations, lowering the operating voltage for read operations can also help reducing the energy consumption

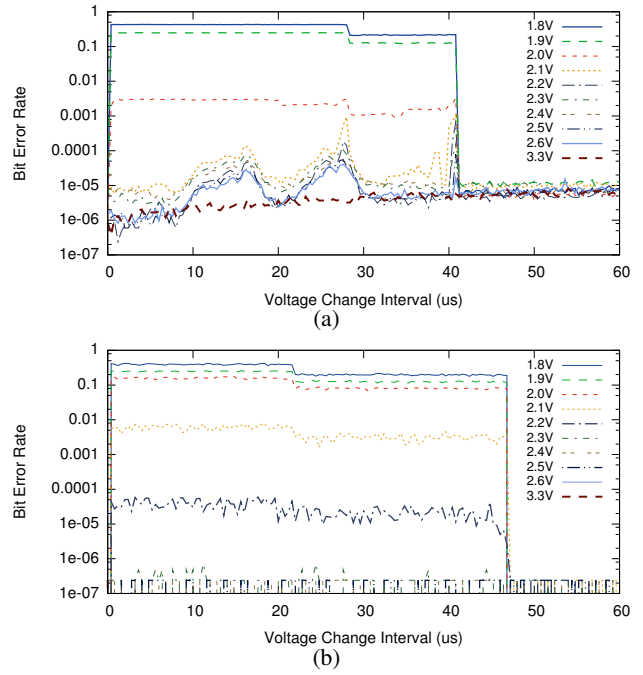


Figure 5: The bit error rate of reading random data from (a) B-MLC128 (b) E-MLC8 when supply voltage droops to 1.8 V - 2.6 V at different voltage change intervals




Figure 6: The latency of reading random data from E-MLC8 when supply voltage droops to 1.8 V - 2.6 V at different voltage change intervals

even though the latency increases. For E-MLC8, we can reduce 36% read energy by operating the chip at 2.5 V.

4.3 Erase and power fade

To examine the effect of power fade for erase operations, we first program a block with random data, and then change the supply voltage while erasing the block. We use the number of zero bits as the bit error rate for erase operations since a block should contain 1s after a successful erase.

The behaviors of erase operations under power fade are similar across all the chips we examined. In general, if the voltage is lower than a certain level, the chip cannot erase all the bits within a block. However, for the supply voltage above a certain level, the chip seems to complete the erase operation.

Figure 7 shows the result of E-MLC8. For 1.8 V, if the voltage change interval is less than $137 \mu s$, the erasing block have about 50% bit error rates, which implies that the erase operation fails completely. For 1.9 V, if the voltage changes at $4.8 \mu s$, the error rate still reaches 50%. However, when the voltage only drops




Figure 7: The bit error rate when erasing E-MLC8 with various voltage change intervals




Figure 8: The program error rates of E-MLC8 of a block when supply voltage droops to 1.8 V – 2.6 V during the previous erase operation

to 2.0 V during the erase operation, the bit error rate reaches 0% (0.1 V lower than voltage required for programming), the operation succeeds.

In [14], the authors show that erasing a block with power failure can increase the error rate of later program operations even if all bits in the block seems to be in the erased state. To examine this problem, we first erase a block of random data with voltage droops. Then, we program random data into the block and measure the program error rate. However, we did not find the same result as power failure. The program error rate shows no difference if the block seems to be erased under a lower voltage level.

Figure 8 uses E-MLC8 to illustrate the result. The x-axis is the voltage change interval for the previous erase operation, and the y-axis is the bit error rate of the later program operations on the same block. The program error rate only becomes high under erasing voltage at 1.9 V and 2.0 V, where erasing the block will have high error rates. When we program a block erased under 2.0 V in E-MLC8, the error rate shows no difference comparing with programming a block erased with 3.3 V.

Figure 9 shows the latency of erase operations of E-MLC8 when we change the supply voltage during an operation. The trend is consistent with other operations. Lowering the supply voltage increases latency. In terms of energy consumption, erasing a block in E-MLC8 at 2.1 V can help reduce the energy consumption by 38%.

5. DYNAMIC VOLTAGE SCALING FOR FLASH MEMORY

As flash memory makes inroads into the mainstream storage devices for embedded systems and sensor networks, the energy consumption of flash memory becomes a design issue in these energy-limited devices. In this section, we will present a dynamic voltage scaling scheme that helps saving energy based on our characterization results.




Figure 9: The latency of erasing E-MLC8 when supply voltage droops to 1.8 V – 2.6 V at different voltage change intervals

Name	Program	Read	Erase
A-MLC16	2.3V	2.5V	2.1V
B-MLC32	2.2V	2.7V	2.0V
B-MLC128	2.1V	2.7V	1.7V
E-MLC8	2.1V	2.5V	2.0V
F-MLC16	2.4V	2.4V	2.2V

Table 3: The minimum operating voltage for each chip we applied in our dynamic voltage scaling scheme

Our experiments in the previous section indicates that lowering the supply voltage during an operation does not always result in higher error rates. High-resolution power measurement also shows that operating flash memory chips at lower voltages can reduce energy consumption. We define the lowest voltage that allows a chip to complete an operation without increasing the error rate using any voltage change interval as the *minimum operating voltage* for the corresponding operation. Table 3 lists the minimum operating voltage for each chip we examined in our work.

To achieve energy saving, our scheme switches the supply voltage to the minimum operating voltage immediately after the chip starts an operation using high-speed transistors. Since the error rates do not increase when the chip is operating at its minimum operating voltage, we do not need to modify the existing error coding scheme. After the operation completes, we immediately switch back to the maximum of minimum operating voltages across all operations for that chip. For example, in E-MLC8, we will switch back to 2.5 V. For B-MLC32 and B-MLC128, because reading at lower than suggested voltage range cannot guarantee the data correctness, we only switch the supply voltage to 2.7 V during read operations. We validated the error rates of all the chips that we examined in this work to make sure these chips function correctly at the maximum of minimum operating voltages.

To understand the potential energy saving and performance

Name	Reads	Description
Berkeley-DB Btree	34%	Transactional updates to a B+tree key/value store
Build	94%	Compilation of the Linux 2.6 kernel
Financial	15%	Live OLTP trace for financial transactions [1].
Patch	83%	Applies patches to the Linux kernel from version 2.6.0 to 2.6.29
OLTP	80%	Real-time processing of SQL transactions
Postmark	97%	Models an email server
ptfdb	97%	Palomar Transient Factory database realtime transient sky survey queries
run14	45%	24 hour trace of a software development work station.
WebSearch3	99%	I/O traces from a popular search engine [1].

Table 4: We use traces from nine benchmarks and workloads to evaluate the dynamic voltage scaling for flash memory chips.




Figure 10: The relative active energy consumption after using the dynamic voltage scaling scheme

impact, we implemented the proposed dynamic voltage scaling scheme into a trace-based simulator. We use the timing and power that we measured in Section 4 for simulation. We also include the overhead of switching among different supply voltages that we described in Section 3. Table 4 summarizes the traces we used in our experiments. These traces cover a variety of applications including Internet services and databases. We ran each trace on our simulator with the latency and energy data that we measured in the previous experiments.

Figure 10 shows the normalized energy consumption (not including idle power) of our dynamic voltage scaling scheme. We normalize the energy value with respect to the energy consumption of the system that does not use any energy saving scheme. On average, we can save 22% to 45% energy using the proposed scheme. For B-MLC32 and B-MLC128, because the supply voltages we use to read data are relatively higher than other chips, the energy saving is smaller, especially in traces consist of intensive reads, such as postmark, WebSearch3, and ptfdb.

In terms of performance, the dynamic voltage scaling mechanism only hurts the performance for less than 0.1% in most benchmarks. The only exception is the run14 trace, where we suffer 2% to 19% performance degradation depending on the chip we used.

6. RELATED WORK

Flash manufactures release very little information about the behavior of flash memory chips under adverse operating conditions. Our work compliments previous efforts[14, 3] to quantify chip-level integrity properties that are otherwise unavailable. Boboila and Desnoyers [2] characterize the timing, endurance, and FTL designs for SSDs. However, none of the above works focuses on data integrity issues for power fade.

Though flash-based storage devices consumes relatively low power than conventional storage technologies, the energy efficiency of flash-based storage system is still an issue for devices that require ultra low power storage systems. The flash-based storage system still contribute 14.5% of a sensor node after using a highly optimized file system [9].

Previous efforts to optimize energy and power consumption in SSDs have tried to utilize heterogeneous storage technologies to maximize the energy efficiency for flash-based storage system [8, 10]. Lee and Chang [8] presents a memory allocation strategy that maximizes the energy efficiency for handheld devices. Mathur, Desnoyers, Ganesan, and Shenoy [10] proposes using a surface-mount NAND flash modules to replace the power hungry on-board NOR flash memory. Song, Choi, Cha, and Ha [12] saves energy for multimedia intensive embedded system by improving the flash file system to avoid redundant data compression for write and bypass page caching for reads. Our paper is orthogonal to the above works and can help further reduce the energy consumption of ex-

isting systems.

The programming data can affect energy consumption for program operation. Joo, Cho, Shin, and Chang [6] proposed an energy-aware data compression scheme that minimizes the flash programming energy rather than data size. Salajegheh, Wang, Fu, Jiang, and Learned-Miller [11] presented software-only coding schemes to allow NOR flash chips to program at lower voltage and save energy. However, our work focus not only on the program operation, but also read and erase operations. We also demonstrate the potential of save flash energy without using special coding schemes for any flash operations.

Revising the chip design helps reducing the energy consumption for flash operations. Ishida et al. [4] uses 3D die stacking to allow multiple flash dies to share a common charge pump. Our work only requires the change in micro-controller but still saves energy.

7. CONCLUSION

This paper examines the data integrity issue when unstable or unreliable power supplies underpower flash memory chips. Underpowering flash memory chips can corrupt both programming data and data already stored in the flash memory. Underpowering flash chips can also affect the correctness of reading data and result in incomplete block erase. However, underpowering flash memory chips is not always harmful. We found that flash memory chips can tolerate voltage droops, and the tolerable range is different from operations. We utilize the result to dynamically adjust supply voltages to flash chips and achieve up to 45% energy saving.

8. REFERENCES

- [1] Umass trace repository. <http://traces.cs.umass.edu/index.php/Storage/Storage>.
- [2] S. Boboila and P. Desnoyers. Write endurance in flash drives: measurements and analysis. In *FAST '10: Proceedings of the 8th USENIX conference on File and storage technologies*, pages 9–9, Berkeley, CA, USA, 2010. USENIX Association.
- [3] L. Grupp, A. Caulfield, J. Coburn, S. Swanson, E. Yaakobi, P. Siegel, and J. Wolf. Characterizing flash memory: Anomalies, observations, and applications. In *MICRO-42: 42nd Annual IEEE/ACM International Symposium on Microarchitecture*, pages 24–33, 12 2009.
- [4] K. Ishida, T. Yasufuku, S. Miyamoto, H. Nakai, M. Takamiya, T. Sakurai, and K. Takeuchi. A 1.8v 30nj adaptive program-voltage (20v) generator for 3d-integrated nand flash ssd. In *Solid-State Circuits Conference - Digest of Technical Papers, 2009. ISSCC 2009. IEEE International*, pages 238–239, 239a, feb. 2009.
- [5] JEDEC. Preconditioning of Plastic Surface Mount Devices Prior to Reliability Testing. <http://www.jedec.org/sites/default/files/docs/22a113F.pdf>.
- [6] Y. Joo, Y. Cho, D. Shin, and N. Chang. Energy-aware data compression for multi-level cell (mlc) flash memory. In *Proceedings of the 44th annual Design Automation Conference, DAC '07*, pages 716–719, New York, NY, USA, 2007. ACM.
- [7] T.-S. Jung, Y.-J. Choi, K.-D. Suh, B.-H. Suh, J.-K. Kim, Y.-H. Lim, Y.-N. Koh, J.-W. Park, K.-J. Lee, J.-H. Park, K.-T. Park, J.-R. Kim, J.-H. Yi, and H.-K. Lim. A 117-mm2 3.3-v only 128-mb multilevel nand flash memory for mass storage applications. *IEEE Journal of Solid-State Circuits*, 31(11):1575–1583, Nov. 1996.
- [8] H. G. Lee and N. Chang. Low-energy heterogeneous non-volatile memory systems for mobile systems. *Journal of Low Power Electronics*, 1:52–62, 2005.
- [9] G. Mathur, P. Desnoyers, D. Ganesan, and P. Shenoy. Capsule: an energy-optimized object storage system for memory-constrained sensor devices. In *Proceedings of the 4th international conference on Embedded networked sensor systems, SenSys '06*, pages 195–208, New York, NY, USA, 2006. ACM.
- [10] G. Mathur, P. Desnoyers, D. Ganesan, and P. Shenoy. Ultra-low power data storage for sensor networks. In *Information Processing in Sensor Networks, 2006. IPSN 2006. The Fifth International Conference on*, pages 374–381, 0-0 2006.
- [11] M. Salajegheh, Y. Wang, K. Fu, A. Jiang, and E. Learned-Miller. Exploiting half-wits: smarter storage for low-power devices. In *Proceedings of the 9th USENIX conference on File and storage technologies, FAST '11*, pages 4–4, Berkeley, CA, USA, 2011. USENIX Association.
- [12] H. Song, S. Choi, H. Cha, and R. Ha. Improving energy efficiency for flash memory based embedded applications. *Journal of System Architecture*, 55:15–24, January 2009.
- [13] K. Takeuchi, T. Tanaka, and T. Tanzawa. A multipage cell architecture for high-speed programming multilevel NAND flash memories. *IEEE Journal of Solid-State Circuits*, 33(8):1228–1238, Aug. 1998.
- [14] H.-W. Tseng, L. M. Grupp, and S. Swanson. Understanding the impact of power loss on flash memory. In *DAC 2011: Proceedings of 48th Design Automation Conference*.