Exploiting Data Longevity for Enhancing the Lifetime of Flash-based Storage Class Memory

WONIL CHOI, The Pennsylvania State University

MOHAMMAD ARJOMAND, The Pennsylvania State University

MYOUNGSOO JUNG, Yonsei University

MAHMUT KANDEMIR, The Pennsylvania State University

Storage-class memory (SCM) combines the benefits of a solid-state memory, such as high-performance and robustness, with the archival capabilities and low cost of conventional hard-disk magnetic storage. Among candidate solid-state nonvolatile memory technologies that could potentially be used to construct SCM, flash memory is a well-established technology and have been widely used in commercially available SCM incarnations. Flash-based SCM enables much better tradeoffs between performance, space and power than disk-based systems. However, write endurance is a significant challenge for a flash-based SCM (each act of writing a bit may slightly damage a cell, so one flash cell can be written 10^4 – 10^5 times, depending on the flash technology, before it becomes unusable). This is a well-documented problem and has received a lot of attention by manufactures that are using some combination of write reduction and wear-leveling techniques for achieving longer lifetime. In an effort to improve flash lifetime, first, by quantifying data longevity in an SCM, we show that a majority of the data stored in a solid-state SCM do not require long retention times provided by flash memory (i.e., up to 10 years in modern devices); second, by exploiting retention time relaxation, we propose a novel mechanism, called Dense-SLC (D-SLC), which enables us perform multiple writes into a cell during each erase cycle for lifetime extension; and finally, we discuss the required changes in the flash management software (FTL) in order to use D-SLC mechanism for extending the lifetime of the solid-state part of an SCM. Using an extensive simulation-based analysis of an SLC flash-based SCM, we demonstrate that D-SLC is able to significantly improve device lifetime (between $5.1\times$ and $8.6\times$) with no performance overhead and also very small changes at the FTL software.

CCS Concepts: • **Information systems** → **Flash memory**; • **Mathematics of computing** → Coding theory;

ACM Reference format:

Wonil Choi, Mohammad Arjomand, Myoungsoo Jung, and Mahmut Kandemir. 2017. Exploiting Data Longevity for Enhancing the Lifetime of Flash-based Storage Class Memory. *Proc. ACM Meas. Anal. Comput. Syst.* 1, 1, Article 20 (June 2017), 26 pages. DOI: <http://dx.doi.org/10.1145/3084458>

1 INTRODUCTION

During the last decade, the CPUs have become power constrained, and scaling of the logic devices no longer results in substantial performance improvement of computers. Therefore, it is imperative to consider developing additional ways for performance improvement. For instance, one might target the memory wall problem and consider how to achieve higher overall performance by changing the memory-storage hierarchy. Looking at the

We thank the anonymous reviewers and our shepherd, Abhishek Chandra, for their constructive feedback and insightful suggestions. This work is supported in part by NSF grants 1302557, 1213052, 1439021, 1302225, 1629129, 1526750, and 1629915 and a grant from Intel. Dr. Jung's research is partly supported by NRF 2016R1C1B2015312 and 2015M3C4A7065645, DOE DE-AC02-05CH 11231, and MSIP IITP-2017-0-01015. 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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2017 ACM. 2476-1249/2017/6-ART20 \$15.00

DOI: <http://dx.doi.org/10.1145/3084458>

Proc. ACM Meas. Anal. Comput. Syst., Vol. 1, No. 1, Article 20. Publication date: June 2017.

conventional memory-storage hierarchy, we observe that there is a large performance-cost gap between DRAM (located near processor) and HDD, and this gap has become larger with the recent technology advances. Bridging this gap has the potential to boost system performance in all kinds of computing systems. This is possible with a high-performance, high-density and low-cost non-volatile memory (NVM) technology whose access time as well as cost-per-bit fall between DRAM and HDD, and is called Storage Class Memory (SCM) [4, 10]. Despite the recent advances in NVM technologies (such as Phase Change Memory [30] and Magnetic RAM [33]), it is quite unlikely to exploit them in SCM in any near future because of their high fabrication costs. Instead, this paper assumes a NAND flash-based SCM which has been widely used in various kinds of commercial systems ranging from laptops and desktops to enterprise computers.

Flash memory stores binary data in the form of a charge, i.e., the amount of electrons it holds. There are two types of popular flash memories: Single-Level Cell (SLC) and Multi-Level Cell (MLC). An SLC flash cell has two voltage states used for storing one-bit information, while an MLC flash cell has more than two states and stores 2 or more bits at each time. SLC is fast and has a long lifetime, but MLC trades off these metrics for higher density. In order to have the benefits of both the technologies in the same system, a flash-based SCM typically has a hierarchical internal structure: there is an SLC Solid State Drive (SSD) [2, 9, 12, 18] with tens of gigabytes capacity at the upper level, and an MLC SSD with terabyte capacity at the lower level. Write endurance is a significant challenge for the SLC SSD in this setup. The reason is that the SLC SSD services a great portion of the incoming traffic which poses high write pressure on it (flash memory suffers from low cell endurance, i.e., each cell can tolerate 10^4 – 10^5 program/erase cycles).

In this paper, we target the lifetime problem of SLC SSD in an SCM and discuss the opportunity for improving it by relaxing the *retention time* of the flash, i.e., the period of time that a flash cell can correctly hold the stored data. The flash devices traditionally have long retention times and are expected to retain data for one or more years. Although this long-term non-volatility is a must for many flash memory applications, there are also some cases where the stored data does *not* require it. For example, within a memory-storage hierarchy, we expect the SLC SSD to handle the I/O requests with short-term longevity, while the I/O requests with long-term longevity are normally handled by the underlying MLC SSD or HDD. Our characterization study in this work confirms this behavior – we measure the longevity of written data into an SLC SSD for a wide range of enterprise workloads taken from the MSR Cambridge I/O suite [27]. We observe that a majority of written data into the SLC SSD of an SCM for all evaluated workloads exhibits very short longevity, i.e., about 90% of written data in these workloads have a longevity of up to 10 hours (it is less than 3 minutes for some applications and less than 10 hours for some others).

The retention time relaxation for flash memory is previously studied by some works [21, 28]. They have shown that, by controlling the amount of charges in a cell during the write process, it is possible to reduce its retention time (more details on the theory behind this property are given in Section 3.2). The prior works mostly use this property for performance improvement of the flash by reducing its write execution time. In this paper, however, we use the retention time relaxation of flash to enhance its lifetime. The main idea is that, by relaxing the retention time of an SLC device, we can have more than two states in a cell. At each given time, similar to the baseline SLC, we use every two states to write one bit information. In this way, a device stores multiple bits (one bit at each time) before it needs an erase, increasing the number writes to cell during one erase cycle, or simply increasing the PWE¹ of the device beyond the conventional SLC flash (i.e., one). Increasing PWE of a device directly translates into lifetime improvement.

Our proposed flash memory design is called Dense-SLC (D-SLC) and its implementation needs two minor changes at FTL. First, the block allocation algorithm in FTL should be modified to enable having blocks with

¹We use the term “page writes per erase cycle” (PWE) as the maximum number of *logical pages* stored in *one physical page* during *one P/E cycle*.

Fig. 1. The internal architecture of an SSD with SLC NAND flash chips.

different retention times and use them for storing data values of different longevity. Our proposed block allocation algorithm does not require any specific mechanism for retention time estimation. Instead, it uses a simple and yet effective page migration scheme that imposes negligible lifetime and bandwidth overhead. Second, the garbage collection algorithm, which is needed for the flash management, is modified to ease the system-level implementation of writing multiple bits in one erase cycle. These modifications are simple to implement in the FTL and need two-bit metadata information per one block. Using a detailed implementation of D-SLC flash memory in DiskSim simulator [1, 3], we evaluate its lifetime and performance efficiency for a large workload set. Our experimental evaluation confirms that a typical implementation of D-SLC is able to improve SLC SSD's lifetime by up to $8.6\times$ ($6.8\times$, on average) with no degradation in the overall system performance.

2 PRELIMINARIES

2.1 SSD and Flash Memory

Figure 1 illustrates the internal architecture of an SSD that is composed of three components: 1) *Host interface* communicates with the host system, queues the incoming requests, and schedules them for services; 2) The *SSD controller* is responsible for processing I/O requests and managing SSD resources by executing *Flash Translation Layer* (FTL) software; 3) A set of NAND flash memory chips as the storage medium, which are connected to the controller via multiple buses.

NAND flash chip: A flash memory has thousands of *blocks* and each block has hundreds of *pages*. Each page is a row of NAND flash cells. Binary values of a cell is represented by its charge holding property. Flash memory has three main operations: *read*, *program* (write), and *erase*. Page is the unit of a read or a write operation, and reprogramming a cell needs to be preceded by an erase. Erase is performed at unit of a block. Due to the need for erase-before-write operation and high latency of an erase, flash memory usually employs an out-of-place update policy, i.e., when updating a data, the page containing the old data is marked as *invalid*, and the new data is written to an arbitrary clean page. The new page is marked as *valid*.

FTL: The FTL implements some important functionalities for flash memory management. We go over two primary FTL's functionalities in below.

- **Address mapping:** On receiving an I/O request, FTL segments it into multiple pages and maps each page onto flash chips separately. Address mapping for a write request is a two-step process. First, a chip-level allocation strategy [17] determines which chip each page should be mapped to. Then, the page is mapped to one of the blocks and a page inside it (block allocation). For each chip, FTL always keeps one clean block as the active block for new page writes. Within an active block, clean pages are used to program data in a sequential order. Once the clean pages run out in the active block, a new active block is assigned. Finally, the mapping

Fig. 2. (a) The ISPP write process for SLC programming; (b) A demonstration of changes in threshold voltage during ISPP; (c) The case for data error because of threshold voltage drift.

information of each page (i.e., chip number, block number in the chip, and page index in the block) is stored in a *mapping table* which is kept by FTL. On receiving a read request, the FTL looks up the mapping table for finding its physical page location.

- **Garbage collection (GC):** When the number of free pages falls below a specific threshold, the FTL triggers a GC procedure to reclaim the invalid pages and makes some pages clean. When a GC is invoked, the target blocks are selected, their valid pages are moved (written) to other places, and finally the blocks are erased. Due to the page migrations and erase operation, a GC generally takes a long time and consumes significant SSD bandwidth [13, 14, 20].

2.2 SLC-based SSD in an SCM

Flash memory conventionally stores one-bit information in each cell (Single-Level Cell or SLC). However, during the last few years, manufactures leverage the ability to store multiple bits in a single cell – cells in recent products can store 3 bits (called Triple-Level Cell or TLC) before which 2-bit cell (Multiple-Level Cell or MLC) was the norm. The multi-bit capability of a cell is provided by enabling multiple voltage states in it – MLC has four voltage states, whereas TLC has eight voltage states (sometimes called voltage levels). Despite of their low cost per bit, the TLC/MLC flash memories have higher access latencies and lower endurance than the SLC [29].

In order to have the benefits of both technologies in the same system, current SCM designs usually rely on a two-level and hybrid flash-based hierarchy. At upper level, there is a small and fast SLC flash-based SSD (with tens of gigabyte capacity), and a dense MLC flash-based SSD (with few terabyte capacity) is used at lower level. In such an architecture, the SLC-based SSD is responsible for servicing most of the I/O traffic and hence its lifetime is very crucial (because of writes). In this work, we focus on enhancing lifetime of the SLC SSD in SCM. However, the studied characterization and our proposed optimization mechanism is general and can be applied to MLC/TLC SSDs. This is left as the future work.

2.3 SLC Flash Memory

Data in a flash cell is stored in the form of a threshold voltage (V_{th}), i.e., the amount of electrons captured in the cell represents different states. The threshold voltage is formed within a fixed-sized voltage window, bounded by a minimum voltage (V_{min}) and a maximum voltage (V_{max}). For instance, in an SLC cell, the entire voltage window is divided into two non-overlapping ranges (two voltage states S1 and S2 for storing binary values “1” and “0”, respectively), separated by a large gap and one reference voltage (V_{ref}), as shown in Figure 1.

Write operation: When the written data is “1”, no action is needed as the cell is initially in the no-charging state or erase state (i.e., State S1 in Figure 1). On writing “0”, the flash memory employs a specific scheme called Incremental Step Pulse Programming (ISPP) [32]. As shown in Figures 2a and 2b, the ISPP applies a sequence of voltage pulses with a fixed duration (T_{pulse}) and staircase-up amplitude (V_{ISPP}) to the cell, in order to form the

desired threshold voltage (V_{target}). After triggering each pulse, the cell state is verified to check if the programmed threshold voltage reaches V_{target} . This process is repeated until the desired voltage is reached. The program time (T_{PROG}) is a proportional to the number of ISPP loops, that is inversely proportional to V_{ISPP} , and can be expressed as follows [15]:

$$T_{PROG} \propto \frac{V_{target} - V_{start}}{V_{ISPP}}$$

Under a fixed V_{ISPP} , the higher the target voltage (V_{target}) is, the longer the program time is.

Read operation: Reading an SLC flash is realized by applying a reference voltage (V_{ref}) and inferring the threshold voltage (V_{th}). If the threshold voltage is larger than the reference voltage ($V_{th} \leq V_{ref}$), the cell state is S1 and its value is “1”; otherwise, the cell state is S2 and its value is “0”. The flash read time is a proportional to the number of voltage sensing/comparisons. Thus, reading from SLC is very fast since it needs only one sensing/comparison.

Errors in SLC flash: Right after a cell is programmed as “0”, the threshold voltage is around the target voltage (V_{target}). However, as time goes by, due to the charge loss, the threshold voltage in the cell drifts and it will finally overlap with the neighboring voltage state. As a result, the cell data is interpreted as “1” when it is read. We call this data corruption *retention error* [7, 8, 24], which is illustrated in Figure 2c. In this error model, the lower tail of the state S2 overlaps the part of the state S1 after a specific elapse time, called *retention time*. In order to avoid fast data corruption and provide years of retention time in current flash products, the target voltage (V_{target}) of the state S2 is conservatively formed to be far from the erase state (S1).

As a flash block experiences more and more erases (or P/E cycles), the voltage drift (charge loss) accelerates. To enable data integrity for long time, vendors specify a guaranteed retention time (e.g., 10 years) and endurance limit (e.g., 100K P/E cycles) for their flash products.

3 RETENTION RELAXATION FOR LIFETIME ENHANCEMENT

SLC flash products normally guarantee one long-term retention duration throughout the whole flash lifetime. Such a long-term reliability requirement is critical, when a flash-based SSD is used as a main I/O storage and a replacement for HDDs. However, when employing SSD in the intermediate layers of a storage system (e.g., as the SCM which is the focus of this work), such a long retention time guarantee can be an overkill. Hence, if the retention time guarantee could be *relaxed* under a specific condition, one could have an opportunity to improve other system requirements such as performance and endurance without least concern about the data loss.

Relaxing the guaranteed retention time has been explored in various kinds of non-volatile memories [16, 31]. Some prior works exploited retention relaxation for improving the write performance of flash memories [21, 28]. The principle behind most of these works is to form the threshold voltage less accurately, and by doing so, they would reduce the number of loops in the ISPP process – that would reduce (improve) the device program time. In this work, however, we leverage retention time relaxation for enhancing the lifetime of SLC flash memories in an SCM. To the best of our knowledge, this paper is among the first works that exploit retention time relaxation for lifetime enhancement in SLC SSDs. We believe that our findings give insights to SSD developers for developing highly-reliable flash storage.

In the next three subsections, we introduce our mechanism by answering the following questions:

- (1) *What is the distribution of data longevity values in a flash-based SCM? Do all the data written into an SSD need the long retention time guarantee of flash memory?* (Section 3.1)
- (2) *Is it practically possible to relax the guaranteed retention time of a flash memory? What is the theory behind it?* (Section 3.2)
- (3) *How can we exploit the retention relaxation for improving the lifetime of flash memory? What kind of architectural and software support is required to implement such a relaxation?* (Section 3.3)

3.1 Distribution of Data Longevity in I/O Workloads

Proc. ACM Meas. Anal. Comput. Syst., Vol. 1, No. 1, Article 20. Publication date: June 2017.

Fig. 4. (a) Conventional SLC with PWE of 1; (b), (c) and (d) our D-SLC with a PWE of “3”.

workloads are *write-intensive*, as our interest is to investigate the data longevity and improve the storage lifetime (read traffics do not have an impact on SSD’s lifetime). Section 4.3 characterizes our workloads.

Figure 3 shows the cumulative distribution function (CDF) of data longevity of I/O blocks stored in the SSD. For the I/O blocks written once in a workload, we set their data longevity to the maximum (e.g., 10 years) and assume that we are not allowed to relax the retention time for them. One can observe from this figure that, for all the examined workloads, a large portion of the written data blocks have a short longevity in the range of few minutes, few hours or few days. Specifically, the 95th-percentile of I/O blocks written in *prxy_0* have a longevity of 3 minutes; it is 10 minutes for *proj_0*; 1 hour for *src1_2*; 10 hours for *hm_0*, *mds_0*, *src2_0*, *usr_0*, *web_0* and *wdev_0*; 1 day for *prn_0*, *prn_1* and *web_1*; and 10 days for *wdev_2*.

To sum up, a majority of data blocks (95th percentile) in all our examined workloads are frequently-updated; hence they do not need such a long retention time guarantee (up to 10 years) provided by the commercial SLC flash memories. In contrast, a small fraction of the write data need a retention time larger than 10 days (the percentage varies between 1% to 10% across our workloads). Using these characteristics of our workloads, we will demonstrate in the next section how one can trade off the short retention times for a prolonged storage lifetime.

We note that reading from a flash cell multiple times may affect its voltage level and affect its retention time. However, the probability of data disturbance due to the intensive reads is quite low (e.g., less than 0.01% [6]). Moreover, we observed that our workloads do not exhibit such excess reads on data. Therefore, the read disturbance is not a big issue in a retention-relaxed flash cell and the focus of our work is retention times related to data longevity.

3.2 Retention Time Relaxation for NAND Flash

To relax the retention time in flash memory, we first need to investigate how long the threshold voltage drifts due to the charge loss. Some prior works have shown that the voltage drift of a flash device is affected by multiple parameters including the *initial threshold voltage*, the *current device wear-out level* (in P/E cycles), and the *fabrication technology*. Pan et al. [28] proposed a detailed model of the voltage drift distance (D_{drift}) for NAND flash memory. We simplify this model by considering the critical factors as below and use it throughout this work:

$$D_{drift} = K_{scale} \times N_{PE} \times \ln(1 + T_{RT}), \quad (1)$$

where N_{PE} and T_{RT} are the number of P/E cycles the cell (block) experienced and the retention time in hour, respectively. K_{scale} is a device-specific constant.

Being aware of the voltage drift behavior, we can reduce the “voltage guard” between the two states by shifting the threshold voltage of the program state (S2 in Figure 1) to the left. By doing so, we can decrease the drift distance between the two states and relax the retention time. Figure 4 gives an example. Figures 4a shows the baseline SLC flash, in which there is a large voltage guard between two states and a long retention time (e.g., 10 years) is guaranteed. Figure 4b shows the case where, by shifting the program state (S2) to the left, we could achieve smaller voltage guard between the two states, which results in shorter retention time compared to the baseline. In this figure, the new program state is named as IS-1 (intermediate state).

Forming a program state with lower threshold voltage (like IS-1 in Figure 4b) is easy in flash memory – we need to calibrate ISPP process such that we can program the cell to the new threshold voltage level. The ISPP controller is a programmable circuit inside a flash chip; one can tune the ISPP parameters such as voltage step (V_{ISPP}), pulse duration (T_{pulse}), and target voltage (V_{taregt}), which collectively determines the number of ISPP loops, and in turn, the program latency. Recall the ISPP mechanism described in Section 2.3. In the conventional SLC flash, the voltage step (V_{ISPP} , amplitude difference between two consecutive pulses) is set to a large value, which helps the flash cell reach S2 (of Figure 4a) quickly by reducing the number of ISPP loops. However, if we use this large voltage step for programming a retention-relaxed cell, it is very likely that we jump over the intermediate threshold voltage (IS-1 in Figure 4b). Therefore, *we need to decrease V_{ISPP} and enable a fine-grained threshold voltage jump in the retention-relaxed flash*. Decreasing V_{ISPP} can increase the number of ISPP loops and the write latency. Note that we want to keep the program latency of a retention-relaxed flash memory same as in the conventional SLC memory. Fortunately, compared to the conventional SLC (Figure 4a), *the target voltage (V_{taregt}) decreases in a retention-relaxed flash (Figure 4b)* – this would reduce the number of ISPP loops and the write latency. To sum up, to keep the number of ISPP loops and the write latency in retention-relaxed flash same as in the conventional SLC, we need to calibrate both V_{ISPP} and V_{taregt} , while pulse duration (T_{pulse}) is fixed.

3.3 Trading-off Retention Time for Higher Lifetime of SLC Flash

The discussion in Section 3.1 reveals that, on the one hand, the data longevity of the written data blocks into SSD (as SCM) is mostly limited to few minutes, few hours or few days in the transactional and enterprise applications, i.e., much shorter than 10 years provided by the current flash products. On the other hand, we showed that it is possible to relax the retention time guarantee of flash memory by calibrating the voltage guard and the ISPP parameters (Section 3.2). In this section, we propose a novel mechanism and show how retention time relaxation can be exploited for achieving longer lifetime in flash memories. We start by defining a new metric for lifetime analysis that helps us describe our proposal more clearly.

3.3.1 Page Write per Erase cycle (PWE) Metric. We define the term “page writes per erase cycle” (PWE) as the maximum number of *logical pages* stored in *one physical page* during *one P/E (erase) cycle*. The conventional SLC flash memory stores one bit data in each cell during each erase cycle, and hence, its PWE is one.

If one can write more than one bit during an erase cycle, and hence increase the PWE, the device stores an increased amount of data during its whole lifetime (i.e., 50K P/E for an SLC flash memory in our setting), or in other words, the device lifetime gets improved. The increase in the number of writes in an erase cycle does not accelerate the cell wear-out [19, 26]. That is, the total amount of electrons that go in and out of a cell in an erase cycle determines the cell wear-out. The amount of electrons that pass through a cell in an erase cycle is limited, no matter how many writes are applied. Note also that increasing PWE of an SLC device does not necessarily mean that it stores more than one bit information at each moment – the device is still SLC (single bit storage at each given time); rather, it means that the device does not need to be erased before reprogramming it. Our main objective is to enhance the SLC flash lifetime by increasing its PWE to values higher than one.

3.3.2 Overview of the Proposed Mechanism. Figure 4 shows a high-level view of our proposed design versus the conventional SLC flash memory. The conventional SLC flash cell (shown in Figure 4a) has two states: S1 or the erase state (value “1”), and S2 (value “0”). There is a large *voltage gap* between these two states, which results in a very long retention time (10 years in this example). This cell stores one bit at each time and reprogramming it requires first erasing it. Thus, during each erase cycle, it stores one bit – its PWE is one.

Figure 4b shows the initial state of our proposed SLC flash design. Similar to the conventional design, it has two states: S1 (value “1”) and IS-1 (value “0”). However, the voltage gap between these two states is *small*, and hence the device retention time is *relaxed* to smaller durations (say few minutes, few hours or few days). In contrast to the conventional SLC, in our proposed design, we do *not* need to erase the cell before reprogramming. Instead, when the current values gets invalid, the cell can store the new value by using higher voltage values. For example, as shown in Figure 4c, the new binary states are IS-1 and IS-2, representing the new binary values “1” and “0”, respectively. As before, the cell stores one bit data at each time (similar to the baseline SLC) and also, the voltage distribution of the new binary states (IS-1 and IS-2) is calibrated for short retention time. Repeating this procedure, the device stores one more bit in the cell by programming it into states IS-2 and S2 for binary values “1” and “0”, respectively. As this example demonstrates, by calibrating the voltage states in an SLC device and having two intermediate states IS-1 and IS-2, one can store three bits (one bit at each time) in one cell before erasing it. This increases the PWE of the SLC flash from *one* in conventional design to *three* in this example, which directly translates to a longer device lifetime.

We emphasize two points. First, it is feasible to partially program a cell multiple times during one erase cycle. A few prior works [11, 23] have experimentally demonstrated that we can gradually increase the threshold voltage of a cell by repeating the process of electron injection. Second, achieving higher lifetime is not free in this approach. In fact, one would need to adjust the ISPP parameters to take advantage of the intermediate states – that would increase the complexity of the ISPP controller (even though it is not that much). As discussed in Section 3.2, to keep the write performance (the number of ISPP loops) in our design similar to that of the conventional SLC, we need to decrease both the voltage step (V_{ISPP}) and the target voltage (V_{taregt}), while maintaining the pulse duration in each loop unchanged, with respect to the conventional one.

In short, the proposed mechanism, named **Dense-SLC** or **D-SLC**, archives a longer lifetime compared to the conventional design by exploiting the relaxed retention time. The only potential problem with D-SLC is that it may increase the number of page migrations inside the SSD. Indeed, if the written value at each round has a longevity longer than the device retention time (now relaxed to few minutes for example), we need to move it to another location to avoid data loss. In the following, we describe the required changes at the FTL and SSD controller that help to get most of the potential benefits of D-SLC while avoiding the potential overheads related to unwanted page migrations.

3.3.3 Detailed Design of D-SLC. The D-SLC flash design is highly scalable, i.e., by controlling the ISPP parameters and calibrating distribution of the voltage states, it is possible to increase the number of voltage states in D-SLC and hence enhance its PWE. However, this is not always beneficial since, by increasing the number of states, either a more accurate write mechanism (or finer-grain ISPP) is required or the inter-state voltage gap is reduced. The former increases the controller’s complexity (or write latency if we do not want to keep the D-SLC’s performance similar to the baseline SLC). The latter results in an exponential decrease in device retention time which in turn increases the number unwanted page migrations.

In order to provide sufficient retention time for the majority of I/O blocks while keeping the PWE level of D-SLC high, we make use of the data longevity characterization presented in Section 3.1 and the drift model in Section 3.2 for the threshold voltage calibration in D-SLC. We categorize the I/O blocks of each workload into four groups based on their longevity (or retention time): longevity of a block is either *less than 1 hour*, *between 1 hour and 10 hours*, *between 10 hours and 3 days*, or *more than 3 days*. For the I/O blocks which are only written

Table 1. Classification of I/O blocks in the studied workloads based on their longevity.

Name	Percentage of I/O Blocks (%)			
	$\leq 1\text{Hour}$	1Hour~10Hours	10Hours~3Days	$\geq 3\text{Days}$
hm_0	59.8	33.7	6.4	0.1
prn_0	73.3	21.9	4.8	0
prn_1	59.3	33.3	7.4	0
proj_0	96.7	2.7	0.5	0.1
prxy_0	96.1	3.1	0.7	0.1
mds_0	66.4	29.6	3.6	0.4
src1_2	87.9	7.9	4.1	0.1
src2_0	72.5	23.3	4.0	0.2
stg_0	62.8	35.1	2.0	0.1
usr_0	72.9	21.9	4.8	0.4
web_0	62.7	28.7	8.4	0.2
web_1	48.3	24.0	27.7	0
wdev_0	62.3	33.7	3.4	0.6
wdev_2	23.7	48.8	27.5	0
rsrch_0	79.7	20.3	0	0

Table 2. Block state mode assignment (2, 4, and 8-state mode) for different I/O retention times as a function of the block age (i.e., the number of erases a block experiences).

Retention Time	Block Age (Number of Erases to the Block)				
	0~10K	10~20K	20~30K	30~40K	40~50K
$\leq 1\text{Hour}$	8	8	8	8	8
1~10Hours	8	8	8	4	4
10Hours~3Days	4	4	4	2	2
$\geq 3\text{Days}$	2	2	2	2	2

once in a workload during the examined duration, we assume the maximum longevity and they belong to the last group (i.e., that with longevity larger than 3 days). Table 1 reports the ratio of the I/O blocks belonging to the four retention time categories for each workload.

We determine the voltage threshold distribution in an SLC flash by using the drift model in Section 3.2 with two optimization goals. First, we want to increase the PWE of the SLC flash for each data longevity category in Table 1 during the entire lifetime of the device. Second, we want to keep the performance of our SLC design close to that of the conventional SLC. By assuming a fixed duration for each pulse in ISPP, we determine V_{ISPP} to keep the number of ISPP loops close to that of the baseline SLC. Following these optimization goals, Table 2 reports the number of voltage states used for storing I/O blocks of our four longevity categories during the entire lifetime of the device. Similar to the baseline, the block's endurance limit is 50K P/Es.

In this study, we limit our calculation to three modes for each cell: it is either in the 2-state mode (i.e., exactly same as the conventional SLC), 4-state mode (i.e., shown by the example in Figure 2), or 8-state mode (i.e., it has 6 tightly-arranged intermediate states). The 8-state mode has the shortest retention time and very suitable for storing data values with short longevity (like those with “less than an hour longevity”). The 2-state mode has the longest retention time and suitable for data values with long longevity (like those with “greater than 3 days

Fig. 5. Block allocation examples; DSLC-FTL maintains three different active blocks, and write (update) data is stored in one of these. When retention time of a block expires, an scrubbing mechanism moves the valid pages to a clean block with lower mode (e.g., from a 8-state mode block to a 4-state mode block).

longevity”). The 4-state mode has a moderate retention time and is mostly used for values with “10 hours to 3 days longevity”. One can observe from this table that, as the device wears out, the drift rate increases and we need to decrease the device state to lower levels to avoid (unwanted) migrations. As an example, this behavior happens for I/O blocks with a longevity of “1–10 hours” that are targeted to 8-state mode in early cycles of the device lifetime, but later are targeted to 4-state mode for P/E cycles larger than 30K.

We use the three modes described above for our FTL design and main evaluation results. We later analyze the sensitivity of the D-SLC’s efficiency to different parameters including the number of voltage states.

3.3.4 FTL Design for D-SLC Support. To support D-SLC in an SSD, two changes are required at the FTL – the block allocation algorithm needs to be modified to enable multiple blocks/pages with different modes, and the garbage collection algorithm needs to be redesigned to enable reprogramming a page without erasing that. The new FTL is called DSLC-FTL. We describe the modifications to DSLC-FTL for D-SLC with three modes (2-state, 4-state, and 8-state modes). However, our methodology is general and can be applied to D-SLC with a different mode configuration.

Block Allocation in DSLC-FTL: Due to the limitations of the write and erase operations, all cells in a single page and all pages in a single block have to be in the same mode in D-SLC. Thus, as opposed to the conventional SLCs that have two block types (clean or used) at each time, D-SLC has four block types in a flash chip – each block is either clean (or empty), a 2-state mode, a 4-state mode, or a 8-state mode block. Also, at each time, D-SLC has three active blocks and active write points corresponding to the three state modes it has. Figure 5 shows the block allocation algorithm used in DSLC-FTL. On arrival of a new I/O block, the FTL assumes that it will have a short longevity and maps it to the 8-state mode active block (❶). The heuristic behind this assumption is that, as shown in Table 1, a majority of the written data have “less than one hour” longevity which, irrespective of device wear-out level, is always mapped to 8-state mode based on mode-assignment in Table 2. If this I/O block gets updated in less than an hour, i.e., the retention time of a 8-state mode block, the new update is also allocated in the (current) 8-state mode active block (❷); so we do not change the block mode, as its history admits its short data longevity. Otherwise, on expiration of the block’s retention time, we read its all valid pages and migrate them to the 4-state mode active block (❸); so the controller downgrades mode of these pages because of retention time violation. We call this mechanism *data scrubbing* which is part of our DSLC-FTL. We follow the same procedure for the I/O blocks mapped to 4-state mode: if their updates come before 4-state mode expiration,

Fig. 6. The state diagram of a 4-state mode block. The transition between any two state modes is performed by a GC invocation.

we keep rewriting them in the (current) 4-state mode active block (④); otherwise, on expiration of the block’s retention time, we move its all valid pages to the 2-state mode active block by invoking data scrubbing (⑤). The I/O data in 2-state mode block always remains in this mode (⑥).

This simple heuristic is easy to implement and it needs two minor changes at the FTL metadata.

- (1) FTL needs keep the retention time information at block granularity (instead of page granularity). Indeed, when the first page is allocated in a block, FTL records the clock tic for that block, and periodically monitors it for expiration.
- (2) FTL also needs 2-bit information per each block to indicate its status: “00” for the clean mode, “01” for the 2-state mode, “10” for the 4-state mode, and “11” for the 8-state mode.

We note that this implementation gives the maximum flexibility to DSLC-FTL and allows to write any incoming data into either of the blocks, depending on its longevity, but one can employ a retention time predictor (like the one proposed in [22]) to avoid the data scrubbing cost. However, we found that such mechanism brings a negligible lifetime gain and the data scrubbing in our scheme imposes a very small overhead (see Section 5.4). Accordingly, the current version of D-SLC exploits data scrubbing, instead of a retention-time predictor.

Garbage Collection in DSLC-FTL: We now describe the garbage collection procedure employed for a 4-state mode block, as an example, in D-SLC. This mechanism can be generalized to other modes as well. The diagram in Figure 6 depicts the life-cycle of a 4-state mode block. At any given time, a block can be in one of the four states:

- (1) **Clean:** A block is initially clean or empty. All pages are erased.
- (2) **Round1:** Starting with a clean block, at this state, we write data into the pages of the block in an in-order fashion (i.e., page $i+1$ has to be written after page i). In this state, we use two first states, i.e., the states S1 and IS-1, for writing one bit data in each cell.
- (3) **Round2:** When all the pages in a block are used up in Round1, the block state is changed to Round2 and we store one new page in the target page frame. Again, due to the constraint imposed by the in-order page writes in a block, the next three following actions have to be sequentially applied: (1) All valid pages, programmed in Round1, have to be relocated to elsewhere; (2) We apply dummy write pulses to all page frames to change their voltage states to IS-1 (pseudo-erase state for Round2) – so, all the pages (cells) in the block will have the state IS-1; and (3) We start the second round by using two intermediate states (IS-1 and IS-2). The writes are again performed in an in-order fashion.
- (4) **Round3:** When all the pages in a block are used up in Round2, the block state is changed to Round3 by following a procedure very similar to what described for Round1 to Round2 transition. The only difference is that, during Round3, FTL uses two last states (IS-2 and S2) for writing one new data.

Table 3. Configuration of the baseline SLC-based SSD.

SSD	64GB capacity, eight 8GB SLC NAND flash chips [25]
Flash chip	8192 blocks per chip, 128 pages per block, 8KB pages, block endurance of 50K P/E cycles
Timing parameters	35us for page read, 350us for page write, 1.5ms for block erase, 200MB/sec data transfer rate of the chip.
FTL	GREEDY garbage collection, and round-robin policy for block allocation.

Here are some salient points to keep in mind about this block diagram:

- Dummy write is the process during which all cells in all pages of a block are initialized to the state “1” in Round2. In fact, when the controller decides to change the status of a block from Round1 to Round2, it needs to make sure that all the cells in the block have the state IS-1 (i.e., like erase state for Round2). Implementing dummy write is easy – at the end of Round1, if the content of a cell is “1” (S1), the controller writes into the cell to make its state IS-1; otherwise, i.e. the cell’s content is “0” (IS-1), no action is required. The same procedure (dummy write) applies at the end of Round2, in order to make sure that all cells have the state IS-2 (i.e., like erase state for Round3).
- Changing the block status from Round1 to Round2 and Round2 to Round3 is carried out by garbage collection (GC). This is because we need to move all valid pages to elsewhere, prior to applying our dummy writes. However, in these cases, we do *not* erase the block.
- When all the pages of a block in Round3 are used up, we invoke a normal GC in order to move remaining valid pages and erase the block after page movements (making it ready for Round1 programming).
- D-SLC can work with any GC algorithm available for flash memories. When FTL invokes GC, it chooses one of the already-used blocks, regardless of its current state (i.e., the selected block can be either in Round1, Round2 or Round3). After moving the valid pages of the victim block, FTL applies an erase pulse (if the current state is Round3) or a dummy write (if the current state is Round1 or Round2). So, we do *not* distinguish among the blocks in Round1, Round2 and Round3 during the victim block selection.

4 EVALUATION METHODOLOGY

4.1 Evaluation Framework

We used DiskSim simulator [3] with the SSD extensions by Microsoft [1] to model an SLC-based SSD as an SCM.² This simulator is highly-parametrized and modularized which enables us configure various parameters including the number of flash chips, the flash internal components (i.e., the number of blocks, the number of pages in a block, and page size), and different timing values (i.e., page read and write latencies, block erase time, and data transfer time in/out of the flash chip).

On top of the DiskSim+SSD simulator, we added one function (*data scrubbing*) and modified two existing functions (*block allocation* and *garbage collection*) for D-SLC and its FTL implementation.

- The *data scrubbing* function implements the data scrubbing mechanism (i.e., when the retention time of a block is expired, valid pages in it, if any, are moved to a new block).
- The *block allocation* algorithm is modified to (i) support and maintain multiple active blocks for each flash chip in D-SLC, and (ii) implement the block allocation algorithm in Section 3.3.4.
- The *garbage collection* algorithm is also modified to support our multiple-round GC policy in Section 3.3.4.

²Our analysis and reported results in this paper are based on simulation. As a part of future work, we plan to have a more realistic implementation of D-SLC and DSLC-FTL by using OpenNVM [34].

Table 4. Important characteristics of our workloads.

Name	Description	Write Ratio (%)	Write Size (KB)	Read Size (KB)
hm_0	Hardware monitoring	64.5	7.4	8.3
prn_0	Print server	89.2	22.8	9.7
prn_1	Print server	24.7	22.5	11.7
proj_0	Project directories	87.5	17.8	40.9
prxy_0	Firewall/web proxy	96.9	8.3	4.6
mds_0	Media server	88.1	23.7	7.2
src1_2	Source control	74.6	19.1	32.5
src2_0	Source control	88.7	8.1	7.1
stg_0	Web staging	84.8	24.9	9.2
usr_0	User home directories	59.6	40.9	10.3
web_0	Web/SQL server	70.1	29.9	8.6
web_1	Web/SQL server	45.9	45.9	9.2
wdev_0	Test web server	79.9	12.6	8.2
wdev_2	Test web server	99.9	6.1	8.1
rsrch_0	Research projects	90.7	10.9	8.7

4.2 Configuration of the Baseline System

Table 3 gives the details of the baseline SSD configuration. It is a 64GB SSD with eight 8GB SLC flash chips. The flash memory parameters are taken from a modern Micron device [25] – each chip has 8K blocks, each block has 128 pages and each page is 8KB. The read, write and erase latencies are 35 microseconds, 230 microseconds, and 1.5 milliseconds, respectively. The block endurance is 50K P/E cycles. We also assume that its FTL uses GREEDY algorithm [5] for victim block selection during garbage collection, and the chip-level allocation strategy is static [17].

4.3 Workloads

We use the I/O traces provided in the MSR Cambridge suite [27]. These I/O traces are collected from different transactional and enterprise applications (or different disk volumes in a system running one single application) running multiple consecutive days, which allows us capture the longevity of I/O blocks for long time durations. Among the 36 traces in this benchmark suite, we used 15 traces for our evaluations. Our workloads are listed in Table 4 (different indices refer to different volumes of the same application). The 21 excluded traces are either read-intensive (their write ratios are less than 20%) where lifetime of the baseline SSD is not a concern (the endurance enhancement is the main goal of our technique), or many blocks in them are accessed once during the trace collection time (that is one week in these traces). Table 4 gives the important characteristics of the studied workloads in terms of the *write ratio*, *average write request size*, and *average read request size* (note that Table 1 reports the retention time categorization of data blocks in these workloads).

4.4 Evaluated Systems

We evaluated and compared the results of three systems:

- (1) **Baseline:** This uses the conventional 2-state mode for all blocks.
- (2) **D-SLC:** This uses the proposed D-SLC technique with three state modes for the blocks (i.e., 2-state mode, 4-state mode, or 8-state mode). It also uses DSLC-FTL for flash management, which includes necessary

Fig. 7. Lifetime analysis of D-SLC and Oracle-D-SLC systems. The results are normalized to the baseline SLC.

functions to implement the scrubbing mechanism, block allocation and garbage collection in D-SLC. As explained before, we assume that D-SLC's read and write latencies in all the block modes are comparable to those in the baseline SLC (and hence there is no latency overhead or enhancement in this design).

- (3) **Oracle-D-SLC:** This system is mostly similar to the D-SLC and uses all DSLC-FTL functionalities except its block allocation algorithm. For block allocation, this system assumes that retention time information for each incoming I/O block is known ahead and hence it is allocated to the most suitable block mode. Doing this, Oracle-D-SLC removes the need for scrubbing and, as a result, gives the upper limit of the lifetime and performance improvement by D-SLC.

During our analysis, the results of the evaluated systems are normalized to the *baseline* system for comparison.

4.5 Metrics

We use the following metrics for our evaluation:

- (1) **Lifetime:** It refers to the lifespan of the SLC SSD system and is measured as the total data volume (in KBs) written to it up to the point that its all chips/blocks reach their endurance limit. Under a fixed endurance limit, the more data written to an SSD, the longer lifetime it has.
- (2) **PWE:** Section 3.3.1 defines our PWE metric. Note that the PWE of the conventional SLC is "1" during its entire lifetime. However, the proposed D-SLC results in various PWE values for each block during its lifetime (it can be "1", "3" or "7" for the 2-state, 4-state or 8-state modes, respectively).
- (3) **GC rate and GC cost:** The GC rate refers to the average number of GC invocations in a time unit, and GC cost represents the average execution time of a GC. The higher GC rate and cost result in lower available bandwidth for normal I/O operations.
- (4) **Scrubbing rate and scrubbing cost:** The scrubbing rate indicates how often our data scrubbing mechanism is triggered (i.e., the ratio of blocks on which the data scrubbing is actually triggered, as a fraction of the total number of blocks used). The scrubbing cost is the average number of page migrations required for each scrubbing initiation.
- (5) **Throughput:** It is measured as the amount of data (in KBs) read from or written to the SSD in a time unit.

5 EVALUATION RESULTS

5.1 Lifetime Enhancement

Figure 7 shows the lifetimes of D-SLC and Oracle-D-SLC, *normalized* to the baseline SLC. Compared to the baseline SLC, D-SLC and Oracle-D-SLC increase the lifetime by $6.8\times$ and $6.9\times$, respectively. Exploiting short retention times and employing multiple state modes (i.e., additional 4 and 8 state modes) are quite effective in prolonging the storage lifespan. Specifically, D-SLC allows more and more data to be written in each P/E cycle by

Fig. 8. Percentages of blocks with 2, 4, and 8 states (whose PWEs are 1, 3, and 7, respectively) during the entire device lifetime.

significantly increasing the PWE, which is analyzed in Section 5.2. We want to highlight that D-SLC achieves a lifetime improvement that is very close (only 1.1% less) to that brought by Oracle-D-SLC. This implies that D-SLC does not need frequent data scrubbing invocations. (Section 5.4 provides an analysis on the data scrubbing overheads).

Compared to the other workloads, D-SLC achieves lower lifetime improvements for web_1 and wdev_2 (5.1× and 4.9×, respectively). This is because over 20% of their data have retention times between 10 hours and 3 days

Fig. 9. Analysis on GC overheads of D-SLC: (a) the numbers of GC invocations per one million writes and (b) the average costs of each GC invocation. The results are normalized to the baseline SLC.

(see Table 1) and such data are not placed in blocks with 8-state mode (see Table 2); as a result, these workloads miss the opportunity for further increasing PWE and improving the storage lifetime.

In contrast, *proj_0* exhibits much higher lifetime enhancement than the average. This impressive result is due to two factors: (1) over 96% of its data have retention times below one hour, thus allowing almost all data to be stored in blocks with 8-state mode and maximizing PWE/lifetime. (2) The block allocation scheme in D-SLC could efficiently separate highly-updated data (with short retention times) from data with long retention times in the workload, which reduces its GC rate and cost (note that the lower the GC rate/cost, the higher the lifetime gain). We note that *prxy_0* with a similar distribution of retention times to *proj_0* cannot lead to such a high improvement. This is because the retention times significantly vary during the workload execution; hence, it cannot reduce the GC rate and cost. We discuss this in more detail in Section 5.3.

5.2 PWE Analysis

The lifetime improvements brought by D-SLC originate from the increase in PWE. Figure 8 shows the PWE analysis for the studied workloads. Each figure shows the percentage of flash blocks with 2, 4, and 8 state-modes (whose PWEs are 1, 3, and 7, respectively) during the whole device lifetime. In general, the larger the gray area (PWE=7), the more beneficial our scheme. Note however that, one cannot directly compare the PWE results of different workloads, since their lifetimes in time (x-axis) are all different. We can observe a few common characteristics across the workloads.

- (1) Throughout the storage lifetime, there are always a few blocks with 2-state mode. This is because a few blocks with 2-state mode (including one of the active blocks) are reserved to serve write data in need of a long retention guarantee. In addition, such blocks have a tendency to maintain their mode (i.e., 2-state mode), as they are not likely to get invalid and erased.
- (2) As the flash device gets older, the number of blocks whose PWE is “3” dramatically increases. This is because the data whose retention times range from 1 to 10 hours should be stored in blocks with 4-state mode from 30K P/Es, while they could be accommodated in 8-state mode blocks at early ages (see Table 2). In the same context, as the storage gets older, the number of blocks whose PWE is “1” also increases (specifically, for *web_1* and *wdev_2*). In these workloads, data whose retention times are between 10 hours and 3 days need to be placed in 2-state mode blocks instead of blocks with 4-state mode from 30K P/Es.
- (3) The ratios of block with different modes continue to change as time goes by. This indicates that each block frequently changes its mode, when it is erased and allocated as a new active block again. Thus, as the workload being executed moves from one phase to another, the storage can adapt to the change and adjust the ratios of blocks with different modes.

Table 5. Analysis on data scrubbing overheads of D-SLC: the data scrubbing rate (percentage of blocks on which the scrubbing is invoked) and the data scrubbing cost (the number of valid pages to be migrated per a block scrubbing).

Workload		hm_0	prn_0	prn_1	proj_0	prxy_0	mds_0	src1_2
Scrub. Rate (%)		0.068	0.012	0.029	0.061	0.187	0.096	0.078
Scrub. Cost (cnt.)		26.01	25.56	26.12	27.88	20.74	16.51	12.89
src2_0	stg_0	usr_0	web_0	web_1	wdev_0	wdev_2	rsrch_0	AVG
0.076	0.021	0.045	0.086	0.109	0.063	0.097	0.028	0.071
13.14	23.42	25.45	26.13	21.11	19.19	25.87	33.49	22.83

5.3 GC Analysis

Figure 9 provides the GC frequencies and the GC costs, both of which collectively analyze the GC overhead of D-SLC. Compared to the baseline SLC, D-SLC decreases “the number of GC invocations per one million writes” and “the cost per a GC invocation” by 9.7% and 6.3%, on average. The reduction in GC overheads helps D-SLC bring additional lifetime and bandwidth benefits, even though the effectiveness is not high. This reduction in GC overheads comes from the isolation of hot data (which are frequently updated) from cold data (with long retention times). Note that D-SLC provides multiple active blocks and groups data with similar retention guarantee together in a single block. Hence, when a GC is invoked, victim blocks (which are 8-state mode blocks in most cases) have a tendency to include relatively fewer valid pages, since no data with long retention times is placed in them. As a result, the number of page migrations during GC decreases, and in turn, new/clean pages are not wasted and the GC invocation frequency is lowered.

The significant lifetime improvement in proj_0 results from the largely-reduced GC overheads as well as its high PWEs. Surprisingly, proj_0 drops “the number of GC invocations per one million writes” and “the cost per a GC invocation” by 16% and 20%, respectively. This indirect advance in lifetime helps proj_0 with our scheme achieve a 8.7x of lifetime improvement, which is far beyond 7x when assuming all blocks whose PWE of “7” are used throughout the storage lifespan. One might note that web_1 also experiences a significantly-reduced GC overheads. Unfortunately, this GC benefit does not lead to the high lifetime improvement (i.e., only 5.1x) in web_1. We want to emphasize that the lifetime enhancement in our scheme mainly comes from the increased PWEs and this additional GC overhead reduction is a secondary advantage.

5.4 Scrubbing Overhead Analysis

Table 5 presents the data scrubbing rate and cost, which collectively represent the data scrubbing overhead. The data scrubbing rate (“the percentage of blocks for which the data scrubbing is triggered as a fraction of the total number of allocated blocks”) is quite low (i.e., 0.071%, on average). Furthermore, the data scrubbing cost (“the number of valid pages in a 128 page-block where the data scrubbing is triggered”) is also low (i.e., 22.83 / 128 pages, on average). This low data scrubbing rate is because most of target blocks are already erased when it comes to the deadline and there is no need to act for such blocks. Note that most (page) data in a block are invalidated before the deadline is reached, and such blocks where most pages are invalidated are the best candidates for the GC. Even though the target block is not erased and the data scrubbing is executed, most of its data are already invalidated, which results in low scrubbing costs.

Compared to other workloads, prxy_0 shows a relatively high data scrubbing rate (0.187%), even though its cost is still low. It is because the longevity of its data blocks vary. One can confirm the impact of this high scrubbing rate from Figure 7; the lifetime improvement of prxy_0 is lowered a bit, compared to Oracle-D-SLC which is aware of the longevity of all data in advance. However, in general, the scrubbing overhead is too small to severely hurt the storage lifetime and bandwidth.

Fig. 10. The I/O throughput brought by D-SLC, normalized to the baseline SLC.

One might wonder why workloads like web_1 and wdev_2, where a large portion of data have long longevity, would experience low data scrubbing overheads. For these two workloads, a large fraction of data have 1 hour to 3 days data longevity, and they are written in the 8-state mode block at first. Once they are moved to 4-state or 2-state mode blocks by the scrubbing, the following updates on these data are written in the 4-state or 2-state active blocks, after which there is no more scrubbing activities on these data. Thus, the scrubbing overhead is not significant for these workloads after the state changes happen.

5.5 Performance Analysis

Figure 10 shows the storage throughput results, which are comparable to those of the baseline SLC. Some workloads experience improved throughput, while others lose a bit of their performance; overall, the storage throughput increases by 3% on average. The important parameters that shape the storage throughput in our D-SLC are three:

- **Device read/write latencies:** If device latencies increase, storage throughput decreases and vice versa. However, our scheme provides read/write latencies close to the baseline SLC, as discussed in Section 3.2. So, we assume that device latencies do not affect the throughput in our scheme.
- **Garbage collection overhead:** The higher the GC overhead, the lower the storage throughput. As evaluated in Section 5.3, our scheme reduces the GC overheads a bit; the saved bandwidth in turn helps the storage throughput increase slightly.
- **Data scrubbing overhead:** This additional storage operation consumes storage bandwidth and has a negative impact on storage throughput. However, as discussed in Section 5.4, our scheme does not frequently invoke the data scrubbing, which minimizes the loss of the storage throughput.

Note that our scheme does *not* have an impact on other critical parameters that might affect the storage performance. For example, the degree of storage parallelization (how many I/O requests the storage can process in parallel) and inter-arrival times (how frequently I/O requests are submitted to the storage) remain unchanged under our scheme and evaluation methodology.

6 SENSITIVITY ANALYSIS

The efficiency of the proposed D-SLC design can be influenced by device parameters or configuration setup. To examine this, we performed a series of sensitivity studies.

Table 6. Block state mode assignment (2, 4, or 8-state mode) for different I/O retention times as a function of the block age in Weak / Normal / Strong devices.

Retention Time	Block Age (Number of Erases to the Block)				
	0~10K	10~20K	20~30K	30~40K	40~50K
≤1 Hour	8/8/8	8/8/8	8/8/8	8/8/8	8/8/8
1 ~10 Hours	8/8/8	4/8/8	4/8/8	4/4/8	4/4/8
10 Hours ~ 3 Days	4/4/8	4/4/4	2/4/4	2/2/4	2/2/4
≥ 3 Days	2/2/2	2/2/2	2/2/2	2/2/2	2/2/2

Fig. 11. Lifetime improvement brought by our scheme in Weak, Normal, and Strong devices, normalized to the baseline SLC.

6.1 Different Voltage Drift Distances

6.1.1 Non-Uniform Device Characteristics. The voltage drift distance as a function of the retention time and the P/E cycles (Section 3.2) can be a little longer or shorter, depending on the device characteristics. Specifically, the drift distance is affected by a wide variety of design factors (such as vendors, technology nodes, material-level characteristics), which makes a need to evaluate our scheme in different devices exhibiting varying drift patterns. In addition to the configuration evaluated in Section 5, we employ two more devices by changing the scaling constant (K) of Equation 1. The three evaluated systems in this experiment are as follows:

- **Weak:** In this device, the voltage state drifts longer under the same P/E cycles and retention times. K is set to 5×10^{-4} .
- **Normal:** This is the configuration employed so far (Section 5). K is set to 4×10^{-4} .
- **Strong:** The voltage state in this device drifts shorter under the same P/E cycles and retention times. K is set to 3×10^{-4} .

These three devices have different mappings of state modes to flash blocks for each pair of P/E cycles and retention times, which are listed in Table 6. For example, the *Strong* device can store data whose retention times are between 1 and 10 hours in blocks with 8-state mode at any time (P/E cycle), whereas in the *Weak* device, such data should be placed only in blocks with 4-state mode after 10K P/Es.

6.1.2 Effectiveness of D-SLC in Different Devices. Figure 11 compares the lifetime improvements achieved by D-SLC in three different devices. As can be seen, *D-SLC brings more lifetime improvements in stronger devices than in weaker devices*. This is because in stronger devices, more data with the same retention times can be stored in blocks with higher PWE values. For example, according to Table 6, the Strong device allows data with retention times between 10 hours and 3 days to be stored in 8-state mode blocks before it reaches 10K P/Es, while such data should be stored in 4-state mode blocks until 10K P/Es in the Normal or Weak devices.

Fig. 12. Percentages of blocks with 2, 4, and 8 states, when running *web_1* on Weak, Normal, and Strong devices.

Fig. 13. Percentages of blocks with 2, 4, and 8 states, when running *prn_0* on Weak, Normal, and Strong devices.

The PWE analysis shows how D-SLC increases PWE values as the target device becomes stronger. As an example, Figures 12a, 12b, and 12c show the percentage of blocks with 2, 4, and 8 states, when running *web_1* in Weak, Normal, and Strong devices, respectively. One can see from these figures that the ratios of blocks with 4 states (black area) and with 2 states (red area) gradually decrease, as the drift resistance increases from Weak to Strong devices. Specifically, the Weak device needs blocks with 2 states to store data whose retention times range from 10 hours to 3 days from around 1,000K hours (i.e., 20K P/Es), whereas such data require 2-state mode blocks after around 1,500K hours (i.e., 30K P/Es) in the Normal device. On the other hand, no 2-state mode block is needed to store such data in Strong device.

One might also observe that, in some workloads such as *prn_0* and *proj_0*, the lifetime improvements stay quite low in the Weak device. This phenomenon can be explained by the PWE analysis. Figures 13a, 13b, and 13c show the percentage of blocks with three different states when executing *prn_0* in the Weak, Normal, and Strong devices, respectively. The ratio of blocks with 4 states is very low in the Normal device, and it is almost removed in the Strong device. In contrast, the percentage of 4-state mode blocks largely increases in the Weak device; the black area in Figure 13a appears remarkably. It is because the data whose retention times range 1 to 10 hours need 4-state mode blocks quite early (i.e., after 10K P/Es), while such data need them after 30K P/Es in the Normal device and none of them throughout the lifespan in the Strong device.

Table 7. Block state mode assignment (2, 4, 5, 6, or 8-state mode) for different I/O retention times as a function of the block age in 2-Mode / 3-Mode / 4-Mode / 5-Mode devices.

Retention Time	Block Age (Number of Erases to the Block)				
	0~10K	10~20K	20~30K	30~40K	40~50K
≤1 Hour	8/8/8/8	8/8/8/8	8/8/8/8	8/8/8/8	8/8/8/8
1 ~10 Hours	8/8/8/8	8/8/8/8	8/8/8/8	2/4/5/6	2/4/5/6
10 Hours~3 Days	2/4/5/6	2/4/5/5	2/4/4/4	2/2/2/2	2/2/2/2
≥ 3 Days	2/2/2/2	2/2/2/2	2/2/2/2	2/2/2/2	2/2/2/2

Fig. 14. Lifetime improvement brought by device under varying numbers of states, normalized to the baseline SLC.

6.2 Different Numbers of State Modes

6.2.1 Varying the Granularity of D-SLC. So far, our scheme has employed “three” different modes, which are 2, 4, and 8-state modes. However, it is possible to manage voltage drifts at finer granularities by employing additional modes such as 5 and 6-state modes. To this end, we evaluate D-SLC supporting different numbers of state modes. In particular, we compare the following four systems:

- **2-Mode configuration:** This system has two state modes: 2 and 8-state modes.
- **3-Mode configuration:** This system has three state modes: 2, 4, and 8-state modes. This is the configuration employed so far.
- **4-Mode configuration:** This system has four state modes: 2, 4, 5, and 8-state modes.
- **5-Mode configuration:** This system has five state modes: 2, 4, 5, 6, and 8-state modes.

Table 7 provides the different mappings of state modes to flash blocks for each pair of P/E cycles and retention times in the four evaluated systems.

6.2.2 Effectiveness of D-SLC under Various Configurations. In general, the more state modes, the longer the device lifetime. However, some workloads significantly benefit from increasing the number of modes, while the lifetime gain is negligible in others. Figure 14 shows the lifetime improvement achieved by the 2-Mode, 3-Mode, 4-Mode, and 5-Mode devices in four representative workloads, which are categorized into two groups.

Fig. 15. Percentages of blocks with different states, when running *src1_2* in 2-Mode, 3-Mode, 4-Mode, and 5-Mode devices.

Fig. 16. Percentages of blocks with different states, when running *wdev_2* in 2-Mode, 3-Mode, 4-Mode, and 5-Mode devices.

- **Low-beneficial workloads:** As shown in Figure 14a, *prxy_0* and *src1_2* benefit less (or negligible) from the increasing number of modes. It is because majority retention times in these workloads are below 10 hours, and the difference among the 2, 3, 4, and 5-Mode devices is the assignment of different voltage modes (2, 4, 5, and 6-state mode, respectively) to blocks whose P/Es are between 30K and 50K (see the last two columns of the second row in Table 7).
- **High-beneficial workloads:** For *web_1* and *wdev_2*, increasing the number of supporting modes leads to a significant device lifetime improvement (Figure 14b). These workloads include a large amount of data whose retention times are between 10 hours and 3 days, which can be placed in blocks with finer-granularity state modes such as 6 and 5-state mode in the 4-Mode/5-Mode devices. In contrast, in the 2-Mode/3-Mode devices, such data are accommodated in blocks with 2 and 4-state mode (see the third row of Table 7).

The PWE analysis shown in Figure 15 illustrates why a low-beneficial workload (*src1_2*) cannot fully draw the full potential of increasing the number of modes, whereas Figure 16 illustrates how a high-beneficial workload (*wdev_2*) experiences a significantly-increased lifespan by supporting more modes. In *src1_2*, the cliff at the latter of its lifespan represents that the blocks “where data whose retention times range from 1 to 10 hours are stored” change their modes, when their P/E cycles go beyond 30K. Such data can use blocks with 2 (red), 4 (black), 5 (yellow), and 6 (purple) states in the 2-Mode, 3-Mode, 4-Mode, and 5-Mode devices, respectively. Consequently, these small differences do not result in a significant lifetime gain. In contrast, *wdev_2* includes a lot of data whose retention times are between 10 hours and 3 days, and such data can be stored in blocks with more states (or higher PWEs) in early P/E cycles (i.e., from 0 to 30K) in the 4-Mode and 5-Mode devices. As a result, one can observe from the 5-Mode device (Figure 16d) that 20% of total blocks gradually change their modes (i.e., red, black, yellow, and purple areas) throughout the device lifespan, which results in much higher PWEs, compared to the continuous low and unchanged PWE (i.e., “1”) for the same blocks in the 2-Mode device.

7 RELATED WORK

Retention time relaxation has been considered as an attractive optimization option for flash memories. Prior works exploit this capability for different purposes and design trade-offs. We categorize the related works in two groups:

- (1) **Using retention time relaxation for enhancing write performance** [21, 28]: The flash write latency based on the ISPP [32] is mainly determined the number of ISPP loops (see the equation of Section 2.3), which is a function of (i) the distance between start and target voltages ($V_{taregt} - V_{start}$) and (ii) staircase-up amplitude (V_{ISPP}). In general, the higher V_{taregt} (and the longer the voltage distance) or the lower V_{ISPP} , the larger the number of ISPP loops (and the longer the write latency), and vice versa. The works in this group attempt to reduce the number of ISPP loops (and the write latency) by increasing V_{ISPP} and placing the target threshold voltage less-accurately (based on the retention time relaxation). In contrast, our D-SLC targeting lifetime enhancements tries to keep the number of ISPP loops similar or very close to the baseline SLC by adjusting (reducing) both V_{ISPP} and V_{taregt} .
- (2) **Using retention time relaxation for enhancing flash lifetime** [22]: Similar to D-SLC, WARM [22] optimizes flash lifetime by taking advantage of retention relaxation. However, there are substantial differences between the two approaches. WARM begins with *a retention-relaxed flash memory which employs refresh mechanism* to avoid data loss. Motivated by the high overhead of refresh for hot data (i.e., those with longevity less than the refresh period), they propose an algorithm for hot data detection and design separate pools of hot and cold blocks for the efficient refresh management. In contrast, D-SLC is a generic design and the baseline should not necessarily be a retention-relaxed flash nor a flash with the refresh support. Furthermore, instead of employing an algorithm to estimate the data longevity, D-SLC includes a heuristic mechanism, which is able to put data with similar data longevity history in the same block. More importantly, D-SLC writes multiple bits into a cell during one erase cycle, while WARM allows just a single-bit write in each erase cycle (as the baseline SLC does). In other words, D-SLC improves the lifetime by increasing PWE, whereas WARM (still keeping the PWE one) achieves it by removing unnecessary refreshes. Therefore, D-SLC and WARM can be combined for further lifetime improvement.

8 CONCLUSIONS

Despite the advances in non-volatile memory technologies, flash-based SCMs are still widely used by commercial computing systems, ranging from laptop and desktop to enterprise systems, to hide the performance-cost gap between DRAM and HDD. However, the limited endurance seems to be the main issue for flash-based SCMs, and is the target of our design and optimization in this paper. Specifically, we make three main contributions in this paper: First, by quantifying data longevity in an SCM, we show that a majority of the data stored in a solid-state SCM do not require long retention times provided by flash memory. Second, by relaxing the guaranteed retention times, we propose a novel mechanism, named Dense-SLC (D-SLC), which enables us perform multiple writes into a cell during each erase cycle for lifetime extensions. Third, we discuss the required changes in the FTL in order to exploit these characteristics for extending the lifetime of solid-state part of an SCM. Using an extensive simulation-based analysis of a flash-based SCM, we demonstrate that our proposed D-SLC is able to significantly improve device lifetime (between $5.1\times$ and $8.6\times$) with no performance overhead and also very small changes in the FTL software.

REFERENCES

- [1] Nitin Agrawal, Vijayan Prabhakaran, Ted Wobber, John D. Davis, Mark Manasse, and Rina Panigrahy. 2008. Design Tradeoffs for SSD Performance. In *USENIX ATC*. <http://www.microsoft.com/en-us/download/details.aspx?id=52332>.
- [2] Frank Berry. 2015. 2015 Enterprise Flash Storage. In *Flash Memory Summit*.

- [3] John S. Bucy, Jiri Schindler, Steven W. Schlosser, and Gregory R. Ganger. 2008. The DiskSim Simulation Environment Version 4.0 Reference Manual. In *CMU-PDL-08-101*. <http://www.pdl.cmu.edu/DiskSim/>.
- [4] G. W. Burr, B. N. Kurdi, J. C. Scott, C. H. Lam, K. Gopalakrishnan, and R. S. Shenoy. 2008. Overview of Candidate Device Technologies for Storage-Class Memory. In *IBM Journal of Research and Development*, Vol. 52, Issue. 4.5.
- [5] Werner Bux and Ilias Iliadis. 2010. Performance of Greedy Garbage Collection in Flash-based Solid-State Drives. In *Journal of Performance Evaluation*, VOL. 67, Issue. 11.
- [6] Yu Cai, Yixin Luo, Saugata Ghose, Erich F. Haratsch, Ken Mai, and Onur Mutlu. 2015. Read Disturb Errors in MLC NAND Flash Memory: Characterization, Mitigation, and Recovery. In *DSN*.
- [7] Yu Cai, Yixin Luo, Erich F. Haratsch, Ken Mai, and Onur Mutlu. 2015. Data Retention in MLC NAND Flash Memory: Characterization, Optimization and Recovery. In *HPCA*.
- [8] Yu Cai, Gulay Yalcin, Onur Mutlu, Erich F. Haratsch, Osman Unsal, Adrian Cristal, and Ken Mai. 2014. Neighbor-Cell Assisted Error Correction for MLC NAND Flash Memories. In *SIGMETRICS*.
- [9] Feng Chen, David A. Koufaty, and Xiaodong Zhang. 2009. Understanding Intrinsic Characteristics and System Implications of Flash Memory Based Solid State Drives. In *SIGMETRICS*.
- [10] R. F. Freitas and W. W. Wilcke. 2008. Storage-Class Memory: The Next Storage System Technology. In *IBM Journal of Research and Development*, Vol. 52, Issue. 4.5.
- [11] Laura M. Grupp, Adrian M. Caulfield, Joel Coburn, Steven Swanson, Eitan Yaakobi, Paul H. Siegel, and Jack K. Wolf. 2009. Characterizing Flash Memory: Anomalies, Observations, and Applications. In *MICRO*.
- [12] Jim Handy. 2016. Flash Technology: Annual Update. In *Flash Memory Summit*.
- [13] Benny Van Houdt. 2013. A Mean Field Model for a Class of Garbage Collection Algorithms in Flash-based Solid State Drives. In *SIGMETRICS*.
- [14] Benny Van Houdt. 2013. Performance of Garbage Collection Algorithms for Flash-based Solid State Drives with Hot/Cold Data. In *Journal of Performance Evaluation*, VOL. 70, Issue 10.
- [15] Jaeyong Jeong, Sangwook Hahn, Sungjin Lee, and Jihong Kim. 2014. Lifetime Improvement of NAND Flash-based Storage Systems Using Dynamic Program and Erase Scaling. In *USENIX FAST*.
- [16] Adwait Jog, Asit K. Mishra, Cong Xu, Yuan Xie, N. Vijaykrishnan, Ravishankar Iyer, and Chita R. Das. 2012. Cache Revive: Architecting Volatile STT-RAM Caches for Enhanced Performance in CMPs. In *DAC*.
- [17] Myoungsoo Jung and Mahmut Kandemir. 2012. An Evaluation of Different Page Allocation Strategies on High-Speed SSDs. In *USENIX HotStorage*.
- [18] Myoungsoo Jung and Mahmut Kandemir. 2013. Revisiting Widely Held SSD Expectations and Rethinking System-Level Implications. In *SIGMETRICS*.
- [19] Jae-Duk Lee, Jeong-Hyuk Choi, Donggun Park, and Kinam Kim. 2003. Degradation of Tunnel Oxide by FN Current Stress and Its Effects on Data Retention Characteristics of 90-nm NAND Flash Memory. In *IRPS*.
- [20] Yongkun Li, Patrick P. C. Lee, and John C. S. Lui. 2013. Stochastic Modeling of Large-Scale Solid-State Storage Systems: Analysis, Design Tradeoffs and Optimization. In *SIGMETRICS*.
- [21] Ren-Shuo Liu, Chia-Lin Yang, and Wei Wu. 2012. Optimizing NAND Flash-Based SSDs via Retention Relaxation. In *USENIX FAST*.
- [22] Yixin Luo, Yu Cai, Saugata Ghose, Jongmoo Choi, and Onur Mutlu. 2015. WARM: Improving NAND Flash Memory Lifetime with Write-hotness Aware Retention Management. In *MSST*.
- [23] Fabio Margaglia, Gala Yadgar, Eitan Yaakobi, Yue Li, Assaf Schuster, and Andre Brinkmann. 2016. The Devil Is in the Details: Implementing Flash Page Reuse with WOM Codes. In *USENIX FAST*.
- [24] Justin Meza, Qiang Wu, Sanjeev Kumar, and Onur Mutlu. 2015. A Large-Scale Study of Flash Memory Failures in the Field. In *SIGMETRICS*.
- [25] Micron. 2014. NAND Flash Memory MT29F32G08ABAAA, MT29F64G08AFAAA, MT29F128G08A[J/K/M]AAA, MT29F256G08AUAAA, MT29F32G08ABCAB, MT29F64G08AECAB, MT29F128G08A[K/M]CAB, MT29F256G08AUCAB.
- [26] Neal Mielke, Hanmant Belgal, Ivan Kalastirsky, Pranav Kalavade, Andrew Kurtz, Qingru Meng, Nick Righos, and Jie Wu. 2004. Flash EEPROM Threshold Instabilities due to Charge Trapping During Program/Erase Cycling. In *IEEE Transactions on Device and Materials Reliability*, Vol. 4, No. 3.
- [27] Dushyanth Narayanan, Austin Donnelly, and Antony Rowstron. 2008. Write Off-Loading: Practical Power Management for Enterprise Storage. In *USENIX FAST*. <http://iota.snia.org/traces/388>.
- [28] Yangyang Pan, Guiqiang Dong, Qi Wu, and Tong Zhang. 2012. Quasi-Nonvolatile SSD: Trading Flash Memory Nonvolatility to Improve Storage System Performance for Enterprise Applications. In *HPCA*.
- [29] Thomas Parnell. 2016. NAND Flash Basics and Error Characteristics. In *Flash Memory Summit*.
- [30] Moinuddin K. Qureshi, Sudhanva Gurumurthi, and Bipin Rajendran. 2011. *Phase Change Memory: From Devices to Systems*. Morgan & Claypool Publishers.

- [31] Clinton Wills Smullen, Vidyabhushan Mohan, Anurag Nigam, Sudhanva Gurumurthi, and Mircea R. Stan. 2011. Relaxing Non-Volatility for Fast and Energy-Efficient STT-RAM Caches. In *HPCA*.
- [32] KangDeug Suh, Byung-Hoon Suh, Young-Ho Lim, Jin-Ki Kim, Young-Joon Choi, Yong-Nam Koh, Sung-Soo Lee, Suk-Chon Kwon, Byung-Soon Choi, Jin-Sun Yum, Jung-Hyuk Choi, Jang-Rae Kim, and Hyung-Kyu Lim. 1995. A 3.3V 32Mb NAND Flash Memory with Incremental Step Pulse Programming Scheme. In *ISSCC*.
- [33] Yuan Xie. 2011. Modeling, Architecture, and Applications for Emerging Memory Technologies. In *IEEE Design and Test Computers*, Vol. 28, No. 1.
- [34] Jie Zhang, Gieseok Park, David Donofrio, Mustafa Shihab, John Shalf, and Myoungsoo Jung. 2015. OpenNVM: An Open-Sourced FPGA-based NVM Controller for Low Level Memory Characterization. In *ICCD*.

Received January 2017; revised April 2017; accepted May 2017