

# Reliability of NAND-Based SSDs: What Field Studies Tell Us

*This paper presents reliability studies of NAND-based SSDs in production environments, subjected to real workloads and operating conditions.*

By BIANCA SCHROEDER, ARIF MERCHANT, AND RAGHAV LAGISETTY

**ABSTRACT** | Solid-state drives (SSDs) based on NAND flash are making deep inroads into data centers as well as the consumer market. In 2016, manufacturers shipped more than 130 million units totaling around 50 Exabytes of storage capacity. As the amount of data stored on solid state drives keeps increasing, it is important to understand the reliability characteristics of these devices. For a long time, our knowledge about flash reliability was derived from controlled experiments in lab environments under synthetic workloads, often using methods for accelerated testing. However, within the last two years, three large-scale field studies have been published that report on the failure behavior of flash devices in production environments subjected to real workloads and operating conditions. The goal of this paper is to provide an overview of what we have learned about flash reliability in production, and where appropriate contrasting it with prior studies performing controlled experiments.

**KEYWORDS** | Failure; field study; flash technology; production systems; reliability; solid-state drives (SSDs); uncorrectable errors

## I. INTRODUCTION

The popularity of solid-state drives (SSDs) based on NAND flash technology has been growing continuously both for use in consumer devices as well as in data center environments. The number of SSDs shipped in 2016 alone exceeded 130 million units totaling around 50 Exabytes of storage capacity. The advantages of SSDs that drive their increasing market share are clear and include superior performance compared to hard-disk drives (HDDs), in particular for random access workloads, and lower power consumption.

What is much less well understood are the reliability characteristics of flash drives. On the one hand, SSDs have some reliability advantages over HDDs: the lack of moving parts removes concerns about problems such as head crashes, scratches of the media or a failing spindle motor, and makes them more robust in hostile environments. On the other hand, flash-based SSDs also introduce new error modes, most notably because flash cells wear out with use.

As increasing amounts of data are being stored on SSDs, understanding their reliability profile becomes crucial. Until very recently, our understanding of flash reliability was solely based on research studies using controlled experiments with a small number of chips in lab environments under synthetic workloads, often using methods for accelerated testing that put the drive through many cycles to synthetically speed up wearout (e.g., following the JEDEC JESD218 and JESD219 standards [1], [2]).

However, within the two years, three field studies have been published that report on the failure behavior of flash devices in production environments subjected to real workloads and operating conditions. The first examines uncorrectable errors in flash-based SSDs in Facebook's servers [3]. The second paper is our own work and reports on a range of different errors and types of hardware failures in SSDs in Google's data centers [4]. The third paper studies fail-stop failures of SSDs at Microsoft data centers [5].

The goal of this paper is to provide an overview of what we have learned about flash reliability in production, and where appropriate contrasting it with prior studies performing controlled lab experiments and common assumptions. We focus on the following aspects:

- the different types of errors experienced by flash drives and their frequency in the field;
- raw bit error rates (RBERs), how they are affected by factors such as wear-out, age and workload, and their relationship with other types of errors;

Manuscript received January 9, 2017; revised June 26, 2017; accepted July 4, 2017.

Date of current version August 18, 2017. (Corresponding author: Bianca Schroeder.)

B. Schroeder is with the Department of Computer Science, University of Toronto, Toronto, ON M5S 1A1, Canada (e-mail: bianca@cs.toronto.edu).

A. Merchant and R. Lagisetty are with Google Inc., Mountain View, CA 94043 USA.

Digital Object Identifier: 10.1109/JPROC.2017.2735969

- uncorrectable errors, their frequency, and how they are affected by various factors;
- the field characteristics of different types of hardware failures, including block failures, chip failures, and the rates of repair and replacement of drives;
- fail-stop events, their symptoms, and prediction;
- a comparison of the reliability of different flash technologies (MLC, eMLC, SLC drives);
- a comparison of the reliability of flash drives and HDDs.

## II. BACKGROUND

In this section, we briefly describe different sources of errors in flash, as well as techniques flash drives utilize to protect against them, and provide an overview of the data used in the three field studies on flash reliability [3]–[5].

### A. Flash-Specific Sources of Errors

Simply put, a flash cell stores data as a charge that is trapped on a floating gate in a transistor. A program (i.e., write) operation injects electrons to the floating gate and an erase operation removes electrons from the floating gate. Data stored in a flash cell can therefore be corrupted if either electrons are inadvertently trapped in a cell or electrons unintentionally leak from the transistor of a cell. Below we describe at a high level the error mechanisms that have been identified in the literature. For a detailed device-level description of flash error mechanisms we refer to another article in the PROCEEDINGS OF THE IEEE [6] or the book by Micheloni *et al.* [7].

- Retention errors: A cell gradually loses charge over time through leakage current.
- Read disturb errors: Reading a page in a block can (unintentionally) charge other cells in the same block, due to the way reads work in NAND flash.
- Write errors: These errors appear in a freshly written block, as an unintended side-effect when writing (programming) a page in this block.
- Wearout: As flash degrades with repeated program/erase (P/E) cycles, the incidence of all error types above increases.

### B. Other Sources of Data Loss or Corruption in SSDs

Besides problems related to the flash media itself, there are also other sources of data loss or corruption in SSDs. One source of potential issues is bugs in a flash drive's firmware, which can cause data to be overwritten or corrupted. For storage systems using HDDs, a prior field study [8] demonstrates that data corruption is a serious concern. While we are not aware of field studies involving SSDs, there have been known incidents of firmware bugs that can result in data loss or corruption. For example, in 2009, Intel halted shipment of some of their drive models due to a data corruption bug in their

firmware [9], and in 2015, Samsung and Apple had to release SSD firmware updates to fix bugs that can result in data loss or corruption [10], [11].

Another potential source of data loss or corruption is the sudden loss of power during a flash drive's operation, as flash internal metadata cached in device-internal volatile memory might get corrupted and flash program or erase operations might be interrupted before proper completion. While power loss of a device might seem like an avoidable problem, power outages still occur even in data centers of major operators [12], [13]. Recent studies by Zheng *et al.* [14] and Tseng *et al.* [15] show that power loss makes SSDs susceptible to errors. For example, Zheng *et al.* report that SSDs, even those targeted at the enterprise market, can experience a range of failure behaviors when power is cut unexpectedly during operation. These range from bit corruption, to truncated writes, to complete device failure. Other power-related concerns that can severely affect the reliability of NAND flash come from the design of the power supply of the drive as shown by Zambelli *et al.* [16].

### C. Device-Level Protection Against Errors

The drives' first line of defense against bit corruption are error-correcting codes (ECCs) that are stored on the drive along with the data. As long as the number of corrupted bits in a code-word is within the ECC's capability, corrupted bits can be corrected and do not become visible to the application. When the number of corrupted bits is too large to be corrected, the error is uncorrectable and the application receives an error. Flash drives also take some proactive measures to prevent future errors: when the reliability of a block seems to be deteriorating, the drive marks it as bad and removes it from future usage. For example, the drives at Google mark a block bad after it experiences an uncorrectable error or a failed program or erase operation. Many drives also go one step further and identify when an entire chip seems to be failing. Higher end commodity drives contain spare chips, so that they can tolerate a bad chip by remapping its contents to a spare chip. The drives at Google simply remove a bad chip from further usage and continue to operate with reduced capacity. Finally, some drives arrange the drive-internal chips in a RAID-like structure to cope with bursty errors, or page-, block-, or chip-level errors [18].

### D. The Data

Our work reports on results from three different recent field studies, each based on a different data set.

The first, by Meza *et al.* [3], examines the majority of flash-based SSDs in Facebook's server fleet comprising many millions of SSD-days of usage (i.e., the number of days that the drives in the study spent in the field sum up to many million days). The drives are all based on MLC flash and have been deployed in five different hardware platforms, where a platform is defined as a combination of the device capacity used in the system, the host interface technology used, and

the number of SSDs in the system. Table 3 in the Appendix provides an overview of the drives and platforms. Their data include information on uncorrectable errors (which in the context of their work are errors uncorrectable by the SSD, but correctable by the host), the amount of data read from and written to a drive, the number of block erase operations and discarded blocks per drive, as well as information on some external factors, such as temperature and bus power consumption.

The second is our own recent work [4], and is based on data collected over a six-year period in data centers at Google and covers ten drive models, whose key features are summarized in Table 4 in the Appendix. These drives are custom designed high-performance drives, based on commodity flash chips, but use a custom PCIe interface, firmware, and driver. The study includes two generations of drives, where all drives of the same generation use the same device driver and firmware. That means that they also use the same ECCs and the same algorithms for wear leveling. This fact makes this study unique, since differences between different drive models in the same generation can be attributed to differences in the underlying flash, as firmware, device driver, etc., are identical. Since the drives are custom drives with custom logging, there are more detailed data available than what Self-Monitoring, Analysis and Reporting Technology (SMART) provides, including data on correctable and other types of errors (rather than only uncorrectable errors) and hardware failures, such as failed chips. Moreover, this is the only one of the three studies that not only includes multilevel-cell (MLC) drives, but also single-level-cell (SLC) and enterprise MLC (eMLC) drives.

The third is recent work by Naryanan *et al.* [5], who examine over half a million SSDs, which are used by cloud applications, in five large and several smaller data-centers at Microsoft. The drives come from two manufacturers, comprise five drive models, and are all based on MLC flash. The key features of these drives are summarized in Table 5 in the Appendix. The data include SMART monitoring on the drives and various measures of server level workload intensity, as well as information on fail-stop failure, which were derived from trouble tickets. The authors define a fail-stop failure as a drive event that propagates to the corresponding server, causing it to be shut down for external (sometimes physical) intervention or investigation. The device will be replaced or repaired subsequently.

### III. AN OPERATOR'S VIEW OF FLASH RELIABILITY IN THE FIELD


While lab studies often focus on understanding lower level aspects of device reliability, such as the various mechanisms leading to bit corruption, operators in the field are more concerned with a higher level view of device reliability as

it relates to issues that are potentially disruptive to applications and/or data center operations. In combination, the three recent field studies cover three different types of such issues: 1) uncorrectable errors, where a read operation encounters more corrupted bits than the drive-internal ECC can handle; 2) the repair and replacement of drives due to suspected hardware problems, which are problematic from an operator's point of view as both are associated with costs as well as downtime; and 3) fail-stop failures, which are drive issues that propagate to the corresponding server, causing it to be shut down for external intervention or investigation. The three sections below summarize baseline statistics for the occurrence of these events in the field based on the three recent studies [3]–[5].

#### A. Drive Repair and Replacements

Fig. 1 shows the fraction of drives that was replaced within four years of being deployed in the field for ten different drive models at Google. These replacements include only replacements that were due to suspected hardware issues with the drive (rather than the retirement of obsolete models and hardware upgrades). We observe significant differences between models. Typical four-year replacement rates are in the 3%–5% range, but at the high end two drive models experience replacement rates closer to 10%. These rates correspond to annual replacement rates of around 1% for most drives and 2.5% for the highest rates.

Table 1 shows for the same drive models the fraction of drives in each population that was sent to repairs within four years of being deployed in the field. A drive is swapped and enters repairs if it develops issues that require manual intervention by a technician. We observe significant differences in the repair rates between different models. While for most drive models 6%–9% of their population at some point required repairs, there are some drive models,



**Fig. 1.** The percentage of flash drives that are replaced due to suspected hardware problems within the first four years of being deployed in the field based on the Google study.


**Table 1** Overview of Prevalence of Factory Bad Blocks and New Bad Blocks Developing in the Field, and the Fraction of Drives for Each Model That Developed Bad Chips, Entered Repairs, and Were Replaced During the First Four Years in the Field.

Model name	MLC-A	MLC-B	MLC-C	MLC-D	SLC-A	SLC-B	SLC-C	SLC-D	eMLC-A	eMLC-B
Drives w/ bad blocks (%)	31.1	79.3	30.7	32.4	39.0	64.6	91.5	64.0	53.8	61.2
Median # bad block	2	3	2	3	2	2	4	3	2	2
Mean # bad block	772	578	555	312	584	570	451	197	1960	557
Drives w/ fact. bad blocks (%)	99.8	99.9	99.8	99.7	100	97.0	97.9	99.8	99.9	100
Median # fact. bad block	1.01e+03	7.84e+02	9.19e+02	9.77e+02	5.00e+01	3.54e+03	2.49e+03	8.20e+01	5.42e+02	1.71e+03
Mean # fact. bad block	1.02e+03	8.05e+02	9.55e+02	9.94e+02	3.74e+02	3.53e+03	2.55e+03	9.75e+01	5.66e+02	1.76e+03
Drives w/ bad chips (%)	5.6	6.5	6.6	4.2	3.8	2.3	1.2	2.5	1.4	1.6
Drives w/ repair (%)	8.8	17.1	8.5	14.6	9.95	30.8	25.7	8.35	10.9	6.2
MTBRepair (days)	13,262	6,134	12,970	5,464	11,402	2,364	2,659	8,547	8,547	14,492
Drives replaced (%)	4.16	9.82	4.14	6.21	5.02	10.31	5.08	5.55	4.37	3.78

e.g., SLC-B and SLC-C, that enter repairs at significantly higher rates (30% and 26%, respectively). The vast majority (96%) of drives that go to repairs, go there only once in their life. Comparing with the rate of replacement, we observe that for most models less than half as many drives are being replaced as being sent to repairs, implying that at least half of all repairs are successful.

## B. Uncorrectable Errors

Fig. 2 shows the fraction of drives that were affected by uncorrectable errors within the first four years of being deployed in the field, based on the Google study [4]. We observe that uncorrectable errors are common: depending on the drive model between 26% to more than 90% of drives experience at least one uncorrectable error. A closer look at the data reveals that, depending on the model, two to six out of 1000 drive days are affected by uncorrectable errors. Recall that all drives of the same generation use the same device driver and firmware, and also use the same type and strength of ECC, so differences in the rate of uncorrectable errors between any of the MLC and SLC drives (which are all of the same generation) are not due to differences in the ECC.



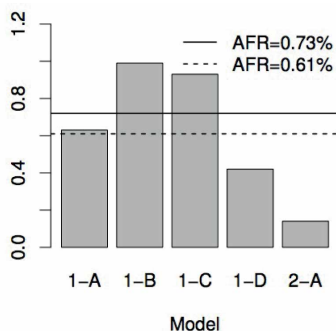
**Fig. 2.** The percentage of flash drives that experience at least one uncorrectable error within the first four years of being deployed in the field based on the Google study.

The Facebook study [3] reports slightly lower incidence rates, ranging from 20% to 35% for most of their hardware platforms, and rates as low as 4%–8% for two of the platforms. The difference might be due to the lower age and less intensive drive usage: the average age of drives range from 0.5 to 2.4 years across their different hardware platforms and drives have on average seen less than 100 P/E cycles, while the Google data spans at least four years for each model and for most models the drives have gone through 500–1000 P/E cycles on average.

The Facebook and Microsoft studies also report uncorrectable bit error rates (UBERs), rather than just error incidence. The range of UBERs across different models in the two studies is very large, spanning  $10^{-11}$  to  $10^{-14}$  at Microsoft and  $10^{-9}$  to  $10^{-11}$  at Facebook. All rates are more than an order of magnitude above the  $10^{-15}$  and  $10^{-16}$  that the JEDEC standard [1] requires for consumer and enterprise class drives, respectively. We speculate that there might be different reasons for those extremely wide ranges of UBER. First, in our own work, we observe that counts of uncorrectable errors have a highly variable distribution with heavy-tails, where a small number of outliers might bias the population mean. Second, we find that UBER might not be a good metric for measuring uncorrectable errors in the field as we will elaborate further in Section V-A. Finally, the Facebook data are based on a nonstandard SSD design, in which the main ECC was performed by the custom driver and the host CPU, and the study defined UBER based on reads that simply required the main ECC engine. It is not clear what the UBER rates are after passing through the host-based ECC engine. The Microsoft data report UBERs after ECCs, but do not mention whether other mechanisms could be applied, beyond ECC, to reduce the host-visible rate of uncorrectable errors.

## C. Fail-Stop Failures

Narayanan et al. [5] focus their study of the Microsoft data on fail-stop failures, which comprise drive events that propagate to the server and cause it to be shut down. According to [5] nearly 80% of the fail-stopped SSDs were replaced. They report annualized fail-stop rates from 0.4% to 1% for four consumer class drive models and around 0.1% for one enterprise class model (see Fig. 3).




**Fig. 3.** The annual percentage of flash drives that experience a fail-stop event based on the Microsoft study (reproduced from [5]).

As the drives in the Microsoft study are commodity drives (unlike Google’s drives, which are custom drives based on commodity chips) it makes sense to compare the fail-stop rates with the annual failure rates one would expect based on vendor specifications. Drive data sheets typically specify a mean time before failure of 1.5–2 million hours, corresponding to an annual failure rate of around 0.4%–0.5%. Narayanan *et al.* quote the specifications for the particular drives in their study in the 0.61%–0.73% range. We note that these numbers from data sheets are consistent with the only field failure rates published by a vendor that we are aware of: Intel reports in a white paper [19] annual failure rates of 0.61%.

Out of the five drive models covered in the Microsoft study, one is barely within the specifications and two exceed the rates in the specifications. What is particularly interesting is that the two drive models whose rates exceed the specifications do not seem to be very heavily used: The amount of data written in relation to drive capacity (see Table 5) implies that the drives on average have undergone fewer than 60–160 P/E cycles, well below the thousands of P/E cycles they should be able to withstand and well before wearout becomes an issue.

As fail-stop failures are closely related to drive repairs and replacements (nearly 80% of affected drives end up being replaced) it is worth comparing with the Google data on repairs and replacements. We observe lower fail-stop rates at Microsoft than drive repairs and replacements at Google. One problem with this comparison is that fail-stop events might not be the only events that trigger drive replacements at Microsoft, so the number of fail-stop events might be lower than actual replacement rates. Other reasons for differences in the numbers could be differences in the usage of the drives. While the drives in the Microsoft study have seen less than 60–160 P/E cycles on average, the drives at Google have on average gone through 500–1000 P/E cycles. Moreover, the Google drive models have larger capacities, ranging from 480 GB to 2 TB, compared to 160–480 GB of the drives at Microsoft.



**Fig. 4.** The most recent published data (figure from [17]) on annual replacement rates for HDDs, collected on systems at Backblaze.

#### D. Comparison With HDDs

An obvious question for operators is how SSDs compare to HDDs, as these two technologies are the main competitors in the persistent storage market.

We observe that replacement rates for SSDs are significantly lower than for HDDs. Work by Schroeder *et al.* [20] and Google [21] shows that annual replacement rates for HDDs typically exceed 1%, with 2%–4% common and up to 13% observed on some systems. These numbers are consistent with the most recent publicly available data on HDDs, reported by Backblaze [17] based on observations on their own data centers (see Fig. 4). In comparison, annual replacement rates for SSDs are significantly lower, as we have seen in Section III-A (typically around 1%, with the worst model at 2.5%).

On the other hand, we observe that SSDs have significantly higher rates of nontransparent errors, i.e., errors that the drive cannot mask from the application. For HDD the main source of nontransparent errors is latent sector errors, where individual sectors on a drive might become unavailable. If a latent sector is discovered during a read access the hard disk has to return an error. Bairavasundaram *et al.* [22] analyze data for 1.5 million hard disks collected from Netapp production storage systems and find that only 3.4% of them develop latent sector errors over a 32-months period (1.5% of enterprise disks and 8.5% of nearline disks). In comparison, for nearly all SSD models, more than 20% of the drives in the field experience an uncorrectable error, and half of all models in the Google study see more than 50% of drives experience uncorrectable errors.

#### IV. RBER IN THE LAB AND IN THE FIELD

A common metric to quantify flash reliability is the RBER, which is computed as the number of corrupted bits divided by the number of bits read (including both correctable and uncorrectable errors).

Table 2 Summary of RBERs for Different Models

Model name	MLC-A	MLC-B	MLC-C	MLC-D	SLC-A	SLC-B	SLC-C	SLC-D	eMLC-A	eMLC-B
Median RBER	2.1e-08	3.2e-08	2.2e-08	2.4e-08	5.4e-09	6.0e-10	5.8e-10	8.5e-09	1.0e-05	2.9e-06
95%ile RBER	2.2e-06	4.6e-07	1.1e-07	1.9e-06	2.8e-07	1.3e-08	3.4e-08	3.3e-08	5.1e-05	2.6e-05
99%ile RBER	5.8e-06	9.1e-07	2.3e-07	2.7e-05	6.2e-06	2.2e-08	3.5e-08	5.3e-08	1.2e-04	4.1e-05

This section is entirely based on the Google study, as the Facebook and Microsoft studies do not contain data on raw bit errors. We can accurately determine the RBER for Google drives in the second generation (i.e., models eMLC-A and eMLC-B), as they report precise counts of the number of bits read and the number of bits corrupted. However, drives in the first generation, while producing accurate counts for the number of bits read, only provide a lower bound on the number of corrupted bits. The reason is that for each page (which consists of 16 data chunks) they only report the number of corrupted bits in the data chunk that had the most corrupted bits. Therefore, in a (not very likely) worst case scenario, where all chunks on a page have the same number of corrupted bits as the worst chunk on the page, the RBER rates would be a factor of 16 larger than the reported numbers indicate. While this does not matter when comparing drives from the same generation, it is a subtlety one must bear in mind when comparing across generations.

### A. RBER Summary Statistics

When studying RBER we find that the average RBER can be dominated by just a few outliers, which can obscure any trends. We therefore report medians and percentiles instead. More precisely, Table 2 reports the median, 95th and 99th percentile of RBER across all drives for a given model.

We observe that RBER varies widely from one drive model to another. For example, the median RBER for drive models in the first generation ranges from as little as  $5.8e-10$  to more than  $3e-08$ . Even drives of the same drive model can have vastly different RBER. For most models, the RBER of a drive in the 99th percentile is about an order of magnitude higher than the RBER of the median drive of the same model.

Finally, we do not find that one vendor produces consistently better RBER than others. For example, among the SLC and eMLC drives, some of the best and some of the worst models come from the same vendor.

### B. What Factors Impact RBER in the Field?

This section studies the impact of the following factors on RBER: wearout from P/E cycles; chronological age, i.e., the amount of time a drive has spent in the field, irrespective of P/E cycles; the number of read, write, and erase operations; and the presence of other errors.

Fig. 5 shows the degree of correlation between RBER and various factors. We used the Spearman rank correlation

coefficient, as it can capture general monotonic relationships, rather than only strictly linear relationships (in contrast to, for example, the Pearson correlation coefficient). We compute the correlation coefficient between the RBER observed for a drive in a given month, and other factors as they were observed in the same month. The factors include the device age in months, the number of prior P/E cycles, the number of read, write or erase operations in that month, the RBER observed in the previous month and the number of uncorrectable errors (UEs) in the previous month. Correlation coefficients close to +1 indicate a strong positive correlation and values close to -1 indicate a strong negative correlation. Each group of bars correspond to one particular factor (see X-axis labels) and each bar within a group corresponds to one drive model. All correlation coefficients are significant at more than 95% confidence.

The figure shows a clear correlation between RBER and all of the factors, except the number of earlier uncorrectable errors, for at least some of the models. However, we stress that some of these correlations might be spurious (as there might be correlations between various factors) and will therefore use the remainder of this section to study the effect of each factor in more detail.

1) *RBER and Wearout*: Due to the limited endurance of flash cells, a drive's RBER rate grows as the number of P/E cycles grows. However, existing studies, based on controlled experiments, do not always agree on the rate of growth. Several papers present graphs that show a superlinear, often exponential, relationship between RBER and P/E cycles [23]–[28]. However, these are all for devices in the 20–30-nm range (compared to mostly 50-nm devices in the


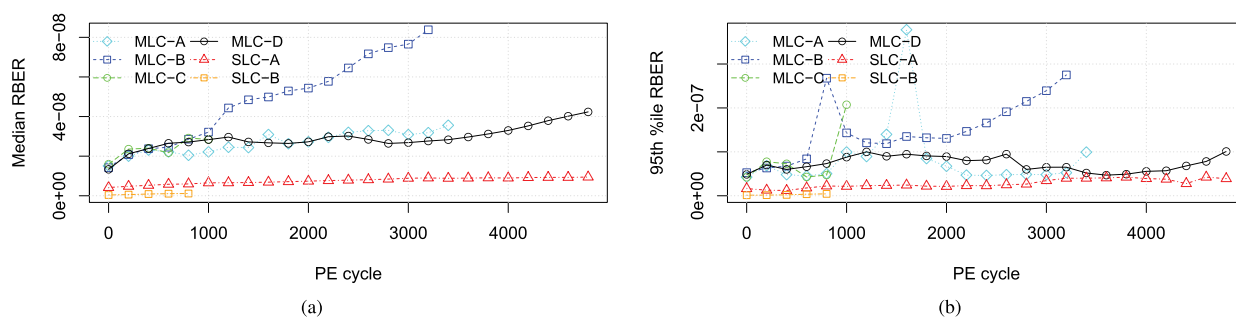


Fig. 5. The Spearman rank correlation coefficient between the RBER observed in a drive month and other factors.



**Fig. 6.** The figures show the median and the 95th percentile RBER as a function of the P/E cycles.

Google data) and for some [23]–[25] the relationship is not clearly visible for smaller P/E cycle counts (less than 10 000), while for another one [29] the rate is exponential only for smaller P/E cycles (less than 5000). Among two papers that report on 50-nm devices, one shows a linear increase for <4000 cycles and superlinear increase for higher cycle counts for some devices [30]. The other paper uses curve fitting and provides an exponential fit to the data [31]. However, visual inspection reveals a poor exponential fit for smaller cycle counts (<4000) and instead indicates a sublinear relationship for this cycle range. Finally, work by Mielke *et al.* [32] shows a sublinear (power-law) relationship for 63–72-nm devices.

Note that all these studies are based on controlled lab experiments and that only one [26] is done at the drive level rather than the chip level. However, we note that the RBER rates reported in this work appear anomalous (rates in the  $1E-11$  range) and might point to problems in the interpretation of the SMART data that the work relied on.

Our goal is to study in detail how RBER grows with P/E cycles in the field. Toward this end, Fig. 6 plots the median and the 95th percentile RBER as a function of the number of P/E cycles. The data points in the two graphs were obtained by sorting all drive days in the Google data into bins, based on the total P/E cycles a drive had seen by this day, and then computing for each bin the median and the 95th percentile RBER.

The figure shows that both median and 95th percentile RBER increase as a function of the number of P/E cycles. We also observe that the growth rate more closely resembles a linear increase, rather than a superlinear or exponential increase as reported in some earlier studies.

We also make another interesting observation: how quickly RBER rates grow with P/E cycles differs greatly between drive models, even for models that start out with comparable RBER rates early in their lifetime. For example, the RBER for all four MLC models is very similar at very low P/E cycles, while there is a 4X difference between the best and the worst model by the time they reach 3000 P/E cycles.

We are also interested in seeing what happens to RBER once a drive goes past its P/E cycle limit and observe that


the growth in RBER continues to be quite smooth (see, for example, model MLC-D with a P/E cycle limit of 3000). Accelerated life tests for the same drives exhibit a steep increase in RBER at around 3X the vendor’s P/E cycle limit. This might indicate that vendor’s P/E cycle limits are chosen very conservatively or that they were based on requirements that we cannot measure with our data. For example, SSD endurance ratings require that a drive maintain a certain number of months (e.g., three months) of data retention capability, so the endurance limit might be chosen at the point where this retention capability is no longer met. It is also possible that the P/E cycle limit was based on early devices of a particular model and that fabrication engineers improved the NAND margins over time.

2) *RBER and Age (Beyond P/E Cycles)*: Lab studies rely on accelerated life tests to evaluate the effect of age on devices. Using field data we can study the effect of natural aging in the field. Here we are particularly interested to see whether there are aging effects besides the P/E cycle induced wear-out of individual flash cells.

Fig. 5 shows a significant correlation between the number of months a drive has been in the field, and its RBER. This correlation might not be surprising, since older drives are more likely to have higher P/E cycles counts, which are correlated with RBER.

We therefore take measures to isolate the effect of age from that of P/E cycles as follows. We assign all drive months to bins based on the deciles of the P/E cycle distribution, e.g., the first bin contains all drive months up to the first decile of the P/E cycle distribution, and so on. Now there is only a negligible correlation between P/E cycles and RBER within a bin as each bin only spans a small P/E cycle range. Next we determine separately for each bin and each drive model the correlation coefficient between RBER and age. Separating by drive model allows us to rule out that any observed correlations are due to differences between younger and older drive models, rather than younger versus older drives within the same model.

We find that even when isolating the effect of the time a drive has spent in field from that of P/E cycles there is still



**Fig. 7. RBER rates as a function of P/E cycles for young and old drives, showing a correlation between age and RBER, independently of P/E cycle induced wearout.**

a significant correlation with RBER (correlation coefficients between 0.2 and 0.4) for all drive models.

Fig. 7 visualizes the correlation between RBER and drive age by plotting RBER as a function of P/E cycles for two separate groups of drive months: The first group consists of data points observed at a young drive age (less than one year) while the second group consists of drive days that were observed when a drive was older (four years or more). The results in Fig. 7 are for one sample drive model (MLC-D). We observe clearly different RBER rates for the two groups, across the entire P/E cycle spectrum.

There could be different reasons for the correlation between age and RBER. One might suspect that the older drives might be from a different (older) manufacturing vintage than younger drives, which leads to a difference in reliability of the two groups; however, we made sure that all drives in the study were put into production within the same one-year time window. One could hypothesize that there are other aging mechanisms at play, beyond cycle-induced wearout, however the literature documents only very few reliability mechanisms that depend simply on age. These typically involve metal interconnects (e.g., corrosion) and would more likely lead to circuit failure rather than individual cell errors. It is also possible that the source of the correlation is related to the drives' workloads. An older drive with the same number of P/E cycles had more retention time between than a younger drive with the same P/E cycles. Longer retention times might have lead to retention errors that increase the RBER. This is an issue that remains to be investigated in more detail in future work.

3) *RBER and Workload*: The workload that a drive experiences can affect its reliability in different ways. As explained in Section II, bit errors can be caused by different mechanisms: retention errors, due to cells leaking charge over time; read disturb errors, where a read operation disturbs the charge in another cell in the same block; or write errors in a recently written block.

While these errors are expected to be correlated with workload, the characteristics of these correlations make them hard to quantify using field data. For example, in the case of read disturb errors, one would need information on

the number of read operations and the RBER per block, as a read can only affect other cells in the same block. One would also need information on the number of read operations between erase cycles, as any erase operations would clear previous errors. One also needs to control for the total number of P/E cycles, when looking at the effect of read operations on RBER, as P/E cycles have an effect by themselves (wearout) and the number of reads and P/E cycles are likely to be correlated. The latter issue can to some degree be controlled by data mining techniques, but the first two issues cannot be easily resolved, as field data is typically not collected at a sufficiently fine granularity (e.g., including per-block information) due to associated overheads. That means when field data shows no correlation between RBER and the number of read operations, we cannot conclude that read disturb does not occur (the per-drive rather than per block data might just be too coarse grained). On the other hand a correlation between RBER and read operations might indicate the presence of read disturb (provided that the analysis controlled for P/E cycles).

The Google study checked for correlations between read operations and RBER and found a correlation (even after controlling for P/E cycles), but only for two of the models (models MLC-B and SLC-B). We conclude that read disturb might affect RBER in the field, at least for some models, but more fine-grained data would be necessary for any definite conclusions.

4) *RBER and Lithography*: To see whether differences in RBER for models using the same technology (MLC or SLC) can be partially explained by differences in feature size, Table 4 includes lithography information for all models in the Google study. As one might expect, models with a smaller lithography tend to have higher RBER. For example, the RBER of the two 34-nm SLC models (models SLC-A and SLC-D) is an order of magnitude higher than that of the two 50-nm models (SLC-B and SLC-C). For the MLC models, the only 43-nm model (MLC-B) has a median RBER that is 50% higher than that of the other three models, which are all 50 nm. Moreover, this difference in RBER increases to 4X with wearout, as shown in Fig. 6.

Differences in lithography might also explain why the RBER for the eMLC drives is several orders of magnitude higher than that of the MLC drives. The two eMLC models are based on 25- and 32-nm chips compared to 50 nm for the MLC drives. (The reader might recall that first generation drives only report a lower bound on their RBER, which in the worst case could be up to 16X higher, however, we observe that that would still result in more than an order of magnitude difference between the eMLC and MLC drives.)

5) *RBER for MLC Versus SLC*: MLC drives are considered to be more susceptible to errors than SLC drives, since MLC cells store multiple bits per cell and as a result the voltage window separating different values is smaller. (See [6] for details on MLC versus SLC.) This difference has been demonstrated in lab studies and is also reflected in the technical



specifications for MLC drives, which have a significantly lower P/E cycle limit.

Revisiting Fig. 6 we observe that there is indeed a significant difference between SLC and MLC drive models, as RBER rates for the MLC models are orders of magnitude higher than for the SLC models. However, we will see in Section V that these differences do not necessarily translate to differences in the rate of uncorrectable errors or other types of errors that are visible to the user.

6) *Effect of Other Factors*: One challenge with field studies compared to controlled lab experiments is that the data available to the data analyst are limited to what the system designers and operators decided to monitor and collect.

While the Google data set is quite rich it cannot account for all possible factors. Examples are details on environmental conditions beyond just average temperatures (e.g., temperature variation), information on external factors, such as power outages, which have been shown to have a detrimental effect on hardware components, or details on workload patterns (beyond just total number of reads and writes), which lab studies have shown to be important.

While we cannot comment on the importance of those individual factors, we do find evidence that existing data cannot directly account for all factors with significant impact on RBER. In particular, we observe that the RBER for a particular drive model varies depending on the cluster where the drive is deployed, even when controlling for P/E cycles.

One example is depicted in Fig. 8, which plots RBER as experienced by MLC-D drives deployed in three different clusters (dashed lines) and compares this to the RBER of the entire MLC-D population (solid line). We verified that these differences cannot be directly explained by any of the factors our data accounts for, including factors, such as physical age, read counts, or write counts.

One possible reason could be different types of workloads running in different clusters. For example, Fig. 8(b) shows that the read/write ratio tends to be higher in those clusters with the highest RBER. However, the read/write ratio does not explain differences across clusters for all models, so it is likely that there are other factors at play, such as other

workload characteristics or environmental factors (temperature, humidity, etc.), that the data do not account for.

The observation above also points out another problem with field studies. Unlike in controlled lab environments, where all factors can be controlled and equalized for all drives under test, many different factors affect the field experience of different drives. Ideally, for the study of any particular effect on a drive in the field one would like to hold all other factors constant, which is not possible for drives in production. When performing our analysis of the Google data, we controlled for as many factors as possible. For example, we ensured that trends we observe persist when studying only drives that are all deployed in the same cluster.

## V. UNCORRECTABLE ERRORS

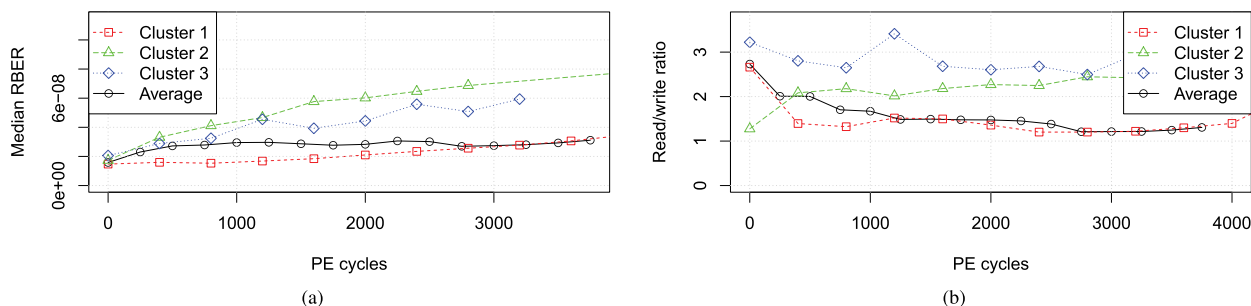
This section takes a closer look at uncorrectable errors (UEs), including a look at metrics to measure UEs and the impact of various factors on UEs.

### A. Why UBER Is Not a Useful Metric for Field Studies

The frequency of uncorrectable errors is commonly measured by a metric called UBER (uncorrectable bit error rate), which is the number of uncorrectable bit errors divided by the total number of bits read. Note that this metric implicitly assumes that the number of uncorrectable bit errors is correlated with the number of bits read, and hence normalizes by this number.

The same assumption underlies the RBER metric and it makes sense in that context, as we find that the number of errors observed in a given month is strongly correlated with the number of reads in the same time period (Spearman correlation coefficient larger than 0.9). The reason for this strong correlation is that one corrupted (but correctable) bit will continue to increment the error count with every read operation that accesses it, since the corrupted bit is not corrected immediately upon detection of the corruption (drives only periodically rewrite pages with corrupted bits).

The assumption also often makes sense in controlled lab experiments, e.g., where the experimenter ensures that



**Fig. 8. (a) Median RBER rates as a function of P/E cycles for model MLC-D for three different clusters. (b) The read/write ratio of the workload for the same model and clusters.**

when different devices are compared they have all experienced the same number of reads and writes, and every bit that was written is being read back. This is, for example, the case if the experimenter follows JEDEC standards.

The same assumption makes less sense for uncorrectable errors in the field: upon the detection of an uncorrectable error the SSD controller will usually remove the corresponding block from further usage, so it will not continue to increase the error count afterwards. To verify this intuition, we employed a variety of measures (Pearson, Kendall, and Spearman correlation coefficients) to study the relationship between the number of reads within a time period and the number of uncorrectable errors in the same time period, as well as visual inspection. Besides the raw number of uncorrectable errors, we also considered the incidence of uncorrectable errors (i.e., the probability that a drive will have at least one within a certain time period) and their correlation with read operations.

None of the measures provides any evidence for a correlation between the number of uncorrectable errors and the number of read operations or the amount of data read. The correlation coefficients are less than 0.02 for all drive models, and visual inspection does not show a higher rate for uncorrectable errors when there are more read operations.


We therefore conclude that UBER is not a meaningful metric to compare the reliability of different drives (or drive types) in the field. UBER is a metric that is more useful in controlled environments where the workload is set by the experimenter and is identical for all drives under test. If used as a metric in the field, UBER will artificially decrease the error rates for drives with high read count and artificially inflate the rates for drives with low read counts, as UEs occur independently of the number of reads.

One might consider an alternative definition of UBER, as suggested in the JESD218 standard, which divides the number of uncorrectable bits by the number of write or erase operations instead of read operations. However, the Google study [4] observes that also write and erase operations are uncorrelated with uncorrectable errors, so such an alternative definition would not be any more useful when considering field data.

## B. What Factors Impact UEs in the Field?

1) *Uncorrectable Errors and Wearout*: Fig. 9 studies how the daily probability of experiencing a UE is affected by a drive's P/E cycles (based on Google data). We note that, similarly to RBER, the probability of UEs grows continuously with P/E cycles and visual inspection as well as curve fitting suggest a linear growth rate.

Also, as was the case for RBER, we observe no significant jump in error probabilities after a drive's P/E cycle limit is exceeded and we see a large variance in error probabilities across models (even those within the same class and



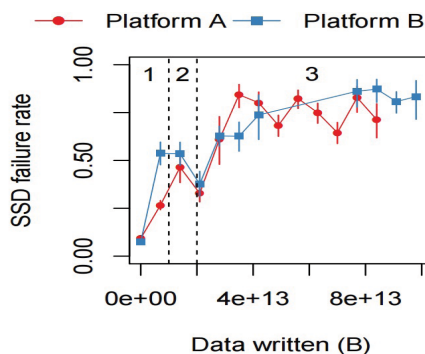
**Fig. 9. Daily probability of a drive experiencing an uncorrectable error as a function of the P/E cycles the drive has experienced.**

similar feature size), albeit the differences are smaller than for RBER.

Interestingly, when comparing Fig. 9 with Fig. 6 on RBER, we observe that the models with the lowest RBER are not necessarily those with the lowest incidence of UEs. For example, at 3000 P/E cycles the RBER rates of MLC-D are 4X lower than those of MLC-B, while its UE probability is actually slightly higher than that of MLC-B. This motivates us to further study the relationship between RBER and uncorrectable errors in Section VIII-B.

2) *Infant Mortality*: Much of the discussion on flash reliability over an SSD's lifetime is focused on wearout with age. However, hardware components are also known to exhibit infant mortality, where many devices fail shortly after being deployed in the field. More generally, a common model for device failure over lifetime is the bathtub model, with high initial failure rates, then lower rates during the useful life, before rates start increasing again when wearout sets in.

Meza et al. [3] have studied the early lifetime behavior of flash drives at Facebook and find that it is slightly more complex than the bathtub model. Rather than initially high rates of infant mortality that drop over time, they observe two distinct periods, as shown in Fig. 10. During the initial period rates increase, then there is a second phase during which rates decrease, before rates start increasing again for the remainder of a drive's life. They call the first phase "early detection period" and hypothesize that during this phase,



**Fig. 10. Rate of uncorrectable errors versus the amount of data written to flash cells during early life. (Reproduced from [3].)**

when the device is first being used in the field, weak cells are being detected and consequently the corresponding block is marked bad and removed from usage. In the second phase reliability improves, since bad cells/blocks have been identified and taken out of circulation. In the third phase, rates start rising again as wearout starts to become the dominating effect.

3) *Uncorrectable Errors and Workload Intensity*: For the same reasons that workload can affect RBER one might expect an effect on UEs. For example, since we observed a correlation between read errors and RBER for two drive models at Google (recall Section IV-B2), read operations might also correlate with uncorrectable errors. Unfortunately, neither the Google data nor the Facebook data exhibit a significant correlation between read operations and the number of UEs (when controlling for P/E cycles). That means we cannot draw any conclusions on the effect of read disturb on UEs, as the lack of a correlation might be due to the coarse-grained nature of the field data used in these studies.

4) *Workload Patterns*: Besides workload intensity, it is also possible that certain workload patterns affect the rate of uncorrectable errors. For example, the amount of time that passes before a previously written piece of data is overwritten or rewritten, will affect the rate of retention errors (which increase as data is left untouched for long periods of time). Unfortunately, there is no field data available that contains low-level input/output (I/O) traces to allow for such analysis (as collecting such data would be associated with considerable overhead).

However Meza *et al.* are able to indirectly study one aspect of workload patterns for the drives at Facebook by observing DRAM buffer usage, which for their devices is used exclusively for drive-internal metadata. They find that as more DRAM buffer is used, the rate of uncorrectable errors increases. They hypothesize that this is because DRAM buffer usage is higher when data is sparsely (e.g., noncontiguously) allocated, as more metadata is needed for the same total amount of data stored. Sparse data allocation might indicate workloads which perform many small writes and hence require a large number of copy and erase operations. The paper unfortunately provides no data on the number of erase operations per drive, which could be used to further validate this hypothesis.

5) *Temperature*: Temperature is well known to affect the reliability of hardware components. This is the reason that for example the JEDEC JESD218 standard requires cycling to be done at both 25 °C and a high-use type of temperature (55 °C–85 °C). Meza *et al.* [3] study the effect of temperature, measured by sensors on the SSD cards in Facebook's fleet, on uncorrectable errors and find that it is not as clear cut as in lab experiments. For the range of operating temperatures observed in their data, spanning 30 °C–65 °C, there was no clear effect in two of the six platforms, for two

platforms errors increased with temperature, and for two platforms errors decreased with temperature. The authors speculate that drive-internal mechanisms deployed by the SSD controller, which try to protect the drive under higher temperatures by throttling workload and power, might explain the stable or decreasing failure rates under higher temperatures for some models.

6) *Lithography*: While we found RBER to be clearly affected by lithography, it is interesting to observe that the effects of lithography are much less obvious in the case of uncorrectable errors. Fig. 9 shows, for example, that SLC-B drives experience a higher rate of uncorrectable errors than SLC-A drives, despite the fact that SLC-B has the larger lithography (50 nm compared to 34 nm for model SLC-A). Moreover, MLC-B, which is the MLC model with the smallest feature size, does not generally have higher rates of uncorrectable errors than the other models. In fact, during the first third of its life (0–1000 P/E cycles) and the last third (>2200 P/E cycles) it has lower rates than, for example, model MLC-D.

While we cannot say for sure why lithography has a weaker effect on uncorrectable errors, it is possible that fabrication process improvements can compensate for the challenges of smaller feature sizes, and that some contributors to UEs like firmware bugs do not depend on lithography at all.

7) *MLC Versus SLC*: eMLC and SLC drives offer a higher write endurance (in terms of the maximum number of P/E cycles they are rated for) than MLC drives and are often expected to be generally more reliable and robust as they are targeted at the enterprise market and command a higher price point. This section uses the Google data to see how accurate this perception is.

A look at Table 2 shows that these expectations are correct with respect to RBER and SLC drives, as RBER is significantly lower for SLC drives than for MLC and eMLC drives. However, we do not find that SLC drives are superior for those reliability metrics that matter most in practice: Neither the rate of repairs and replacements nor the rate of nontransparent errors is lower for SLC drives compared to MLC or eMLC drives (within the P/E cycle ranges that the data covers).

Maybe surprisingly the eMLC drives experience higher RBER than the MLC drives, however recall that their smaller lithography might be responsible, rather than other differences in technology.

We conclude that while SLC drives might be more reliable at very high cycle counts, they are not generally more reliable than MLC drives when comparing the two drive types within the cycle limit of MLC drives.

### C. Correlations Between Errors


This sections looks at whether there are correlations between errors, both within a drive and across drives in the

same machine. Correlations can have practical implications, which make their study worthwhile. First, correlated errors have a higher chance of negatively affecting applications as they are more likely to break redundancy than isolated errors. Second, correlations can indicate some potential to predict errors, as prior errors can be used as an indicator for a higher chance of future errors.

1) *Correlations Between UEs on the Same Drive:* Both the Google and the Facebook studies provide clear evidence of correlations between UEs on the same drive. The Facebook study observes that a large fraction of errors (more than 80%) is concentrated in a small subset of the drive population (10%) and that 99% of drives that have an uncorrectable error have another one within the following week. Some of these correlations might be because the same weak block or cell might generate multiple errors. However, using the Google data we can show that there is correlation beyond this effect: the drives at Google retire a block after it experiences an uncorrectable error, so the same block or cell cannot create multiple errors. We still observe that a drive that has an uncorrectable error has a 30% chance of having another uncorrectable error within the next month, which is significantly higher than the error probability in an average month.

2) *Correlations Between UEs on Different Drives:* The Facebook study also observes that there are correlations between errors in different SSDs on the same machine. They report that an SSD experiencing an uncorrectable error increases the probability of another SSD in the same machine developing UEs by up to 26%. A possible explanation is that SSDs in the same machine share the same or similar operating conditions, such as temperature or workload. Correlations between different SSDs in the same machine are of particular concern in settings where both SSDs are part of the same redundancy group, e.g., part of the same RAID group.

3) *Correlations Between UEs and Other Types of Errors:* The Google data contain information on other types of errors, besides uncorrectable errors, which allows us to investigate whether other types of errors are correlated with uncorrectable errors. To this end Fig. 11 shows the probability of seeing an uncorrectable error in a given drive month depending



**Fig. 11.** The monthly probability of a UE as a function of whether there were previous errors of various types.

on whether the drive saw different types of errors at some previous point in its life (yellow) or in the previous month (green bars) and compares it to the probability of seeing an uncorrectable error in an average month (red bar).

Fig. 11 shows that nearly any type of prior error raises the probability of experiencing uncorrectable errors. The only exception is RBER (prior correctable errors), which is not correlated with later uncorrectable errors. The strongest effect is observed if the prior error was also an uncorrectable error and if the prior error occurred recently (i.e., in the previous month, green bar, versus just at any prior time, yellow bar). For example, while the chance of seeing an uncorrectable error in a random month is only 2%, this probability increases to almost 30% if there was an uncorrectable error in the previous month. Additionally, a number of other types of errors increase the UE probability by more than 5X. For a detailed description of those error types, see [4].

In summary, prior errors, in particular prior uncorrectable errors, increase the chance of later uncorrectable errors by more than an order of magnitude.


## VI. HARDWARE FAILURES

The Google study also includes information on hardware failures, including bad blocks and bad chips. We describe some of the results below.

### A. Bad Blocks

Blocks are the unit at which erase operations are performed. Many drives are shipped with factory bad blocks, i.e. blocks that have already been marked as bad by the manufacturer. Drives can also develop new bad blocks after being deployed in the field.

Drives usually support mechanisms for dealing with bad blocks. Typically, once a block is diagnosed as bad, a drive will try to recover any data that might still be on it, copy it to a different block and stop using the bad block in the future. How blocks are identified as bad is drive policy dependent. For the drives in the Google study, an uncorrectable error, a



**Fig. 12.** The graph shows the median number of bad blocks a drive will develop, as a function of how many bad blocks it has already developed.

write error, or an erase error on a block will lead to the block being marked as bad.

Table 1 summarizes the prevalence of bad blocks based on the Google study. It provides the percentage of drives that developed bad blocks after being deployed, and the median and mean number of bad blocks for drives with bad blocks. The table is based on drives that were deployed four or more years before the time of the study and includes only data for the first four years in the field. The table also includes the corresponding statistics for factory bad blocks.

1) *Bad Blocks Developed in the Field*: Bad blocks are a common occurrence in the field with 30%–80% of drives affected by them. A closer look at the empirical distribution function of the number of bad blocks per drive shows that it is highly variable with a long tail: most drives with bad blocks develop only a small number of them (medians are in the 2–4 range), but once a drive exceeds this number it is likely to develop many more bad blocks. This point is illustrated in Fig. 12. Intuitively, the figure is meant to show how many more bad blocks a drive will likely develop during its lifetime, based on how many bad blocks it has already experienced. More precisely, the figure looks at the statistical distribution of the number of bad blocks (not including factory bad blocks) a drive develops over its lifetime and plots the conditional median. For example, the  $(x, y)$  point says, out of all drives that have developed at least  $x$  bad blocks during their lifetime, how many more than  $x$  did they develop in total. Fig. 12 includes a separate line for each drive model, with MLC models drawn in blue solid lines and SLC models drawn in red dashed lines. We observe, for example for MLC drives a sharp jump in the expected number of bad blocks once a second bad block is detected: 50% of those drives that develop two bad blocks will develop close to 200 or more bad blocks in total.

There is also another interesting interpretation of this data. A drive that experiences hundreds of bad blocks is likely experiencing a chip failure. So our observations above imply that once a drive sees just 2–4 new bad blocks it is very likely to experience a chip failure and that chips seem to either develop only one or two isolated bad blocks, or the chip fails entirely. An interesting direction for future work would be to investigate the potential for predicting chip failure, based on prior bad blocks and possibly other factors (P/E cycles, workload, etc.).

It is also interesting to consider whether bad blocks are typically detected in a way that is user transparent (e.g., in a write or erase operation) or in a way that leads to an uncorrectable error that has to be passed on to the user and can potentially mean data loss.

The Google data do not include detailed records on how each block was identified as failed. However, we observe that the number of erase and write errors is lower than that of uncorrectable errors, which means that bad blocks are most commonly encountered during a read operation making them visible to the user application.

2) *Factory Bad Blocks*: For most drive models, including MLC as well as SLC drives, the vast majority (more than 97%) of drives are shipped with factory bad blocks. The mean and median number of factory bad blocks varies considerably by model. For example, two SLC models see medians of less than 100 factory bad blocks, while most other models see 800 or more factory bad blocks. The distribution of factory bad blocks looks close to a normal distribution, with mean and median being close in value.

Interestingly, the number of factory bad blocks shows a correlation with the number of bad blocks and a few types of errors a drive experiences later in the field. For example, for most models the drives above the 95th percentile of factory bad blocks experience a higher rate of uncorrectable errors in the field, compared to an average drive of the same model. The drives in the bottom 5th percentile are less often affected by timeout errors than an average drive.

## B. Bad Chips

Similar to how bad blocks are being remapped to other blocks, many drives also include mechanisms for dealing with bad chips. For example, higher end commodity drives can include spare flash chips that can be used to replace a chip that dies. The Google drives do not contain spare chips. Instead, when they detect a bad chip, they mark it as failed and continue to operate with reduced capacity. The Google drives mark a chip as bad if the number of errors it experiences within a certain time window exceed some limit or if some predefined percentage of its blocks are bad.

Table 1 shows that failed chips are not a rare occurrence. Depending on the drive model, between 2% and 7% of drives in a population experience bad chips within four years of being deployed in the field. That means that without mechanisms for tolerating failed chips, 2%–7% of drives would have undergone repairs or been returned to the manufacturer.

Recall that Google drives mark a chip as bad either because some percentage (5% for the models in the study) of blocks on the chip are bad or because the chip exceeded some limit on the number of errors it experienced. It is interesting to consider which of these two is the more common reasons for declaring a chip failed, in particular because vendor guarantees for all flash chips in the study state that at most 2% of a chip's blocks will fail within the P/E cycle limit. We find that in two thirds of the cases, a chip was marked bad because of the number of failed blocks it had experienced. That means these are chips that violate vendor data sheets.

## VII. WHAT ARE SYMPTOMS AND PREDICTORS OF FAIL-STOP FAILURES?

Narayanan *et al.* [5] focus their study on the nature of fail-stop failures, which they define as drive events that propagate to the corresponding server, causing it to be

shutdown for external intervention or investigation and consequently lead to drive replacement or repair. They are particularly interested in symptoms and predictors of fail-stop events and identify four drive-level symptoms that significantly increase the likelihood of the drive experiencing a fail-stop event:

- Data errors: These are errors triggered by the cyclic-redundancy-check (CRC) codes and error correcting codes used by the drive to detect and correct data corruption.
- Program or erase failures: These correspond to program or erase operations that failed. They are often symptomatic of block or chip failures.
- SATA downshift: The SATA interface might decide to lower the signaling rate if error counts exceed some threshold.
- Reallocated sectors: The number of sectors that the drive declared bad. These sectors are remapped to other sectors and removed from further usage hence reducing the spare space on the drive (which is also used for other purposes than reallocation, most notably wear leveling).

Fig. 13 shows for one of the drive models in the Microsoft study how the annual percentage of drives experiencing fail-stop events differs between drives that experienced one of the symptoms above, compared to those who did not. The graph shows that data errors have the strongest impact on the occurrence of fail-stop events: drives that experienced data errors have a 20X higher rate of fail-stop events. All other symptoms each increase the rate of fail-stop events by at least 2.75X.

To gauge the potential for predicting fail-stop events based on these symptoms, Narayanan et al. also study what fraction of drives with fail-stop events did have one of these prior symptoms (see Fig. 14). They find that while the majority of fail-stopped drives had at least one of these symptoms (see right-most bar labeled “any”), a significant fraction (38%) did not experience any of these symptoms.

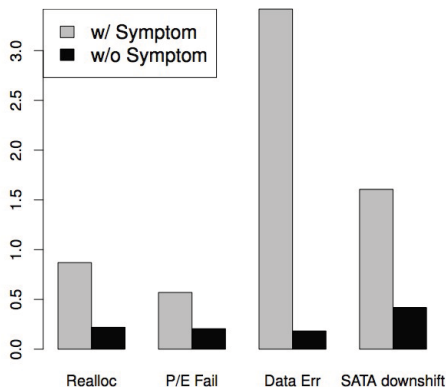


Fig. 13. Annual percentage of drives with fail-stop events in the presence and absence of symptoms (reproduced from [5]).

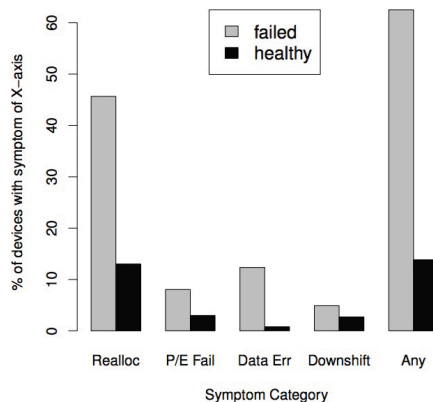


Fig. 14. Percentage of fail-stopped versus healthy drives that exhibit symptoms (reproduced from [5]).

Hence these symptoms are not sufficient in predicting fail-stop events. However, the authors show that taking workload information (most importantly the number of writes) into account can result in more powerful predictors of fail-stop events.


### VIII. FORECASTING FIELD RELIABILITY

Data center operators as well as academics are interested in forecasting how a given device will perform when deployed at large scale in the field, and exposed to real workloads and operating conditions, aging and wearout effects. Typically, two techniques are used: accelerated life tests and projections of the expected rate of uncorrectable errors based on the measured RBER. In this section, we look at how such projections reflect field experience, based on the Google data.

#### A. Accelerated Life Tests

In order to predict the reliability characteristics of a particular device throughout its lifetime, when aging and wearout become a factor, it is common to use techniques for test acceleration, where synthetic workloads are used to add P/E cycles to a drive much faster than they would typically occur in the field. Such techniques are used in many research studies and during the procurement phase in industry. The JEDEC JESD218 and JESD219 provide standards for running such accelerated tests [1], [2].

We obtained data from accelerated tests performed at Google during the procurement phase for some of the devices included in this study. Unfortunately, we have not been able to obtain detailed information on how the tests were run, but as a company with a large deal of experience in procuring and deploying massive amounts of hardware we assume that industry best practices have been followed. A study of this data shows that, for those drive models that



**Fig. 15.** The two figures show the relationship between RBER and uncorrectable errors for different disk models (left) and for individual disk drives within the same model (right).

we had data for, RBER in the field is markedly higher than what the accelerated tests had indicated. As an example, the RBER that eMLC-A drives experienced in the field at an average of 600 P/E cycles, was only reached for more than 4000 P/E cycles in the accelerated tests. These observations serve to show that forecasting field RBER based on accelerated testing is not a trivial exercise.

Another observation we make is that some error mechanisms seem to be difficult to trigger in accelerated testing. In particular, we find that uncorrectable errors and bad blocks were observed only at very high P/E cycles in accelerated tests, while they occur frequently in the field. For example, the six MLC-B drives that were tested during procurement did not experience uncorrectable errors or bad blocks until reaching nearly 10000 P/E cycles (more than three times the limit), yet well above half of these drives developed such problems in the field.

Our study of previous work that publishes RBER based on lab experiments shows a large range of reported RBER numbers, including results that are higher than what we observe in the field. For example, Grupp *et al.* [30], [33] report end-of-life RBER for devices similar to the ones in the Google study (between 25 and 50 nm) in the range of  $1e-08$  to  $1e-03$ , with most common values close to  $1e-06$ . In contrast, the three drive models in the Google study that reach their P/E cycle limit report RBER between  $3e-08$  to  $8e-08$ . Even looking at the 95th percentile of RBER the field rates are significantly lower.

We conclude that predicting field behavior of flash drives based on accelerated life tests is not trivial. One of the main difficulties is likely that workload characteristics in the field can vary widely and are not always captured by standard tests. For example, error rates might be higher under test than in the field because there are mechanisms that allow flash devices to partially recover in the delay between cycles [34] and accelerated cycling therefore provides fewer opportunities to recover between cycles. On the other hand there are also workload-related reasons why error rates in the field can turn out higher than under test. For example, recent work shows how the reading of data before blocks are

fully programmed can increase read disturb errors [35], and hence lab tests, which read only fully programmed blocks, might have lower error rates than field usage, if the field workload frequently reads the most recently written sectors. Also, if the retention time between writes to a data block are longer in the field than the test assumed, field error rates can be higher than error rates under test.

## B. Projecting Drive Reliability Based on RBER

Operators are concerned about raw bit error rates, as high raw bit error rates might turn into uncorrectable errors. Bit errors, as long as they are correctable, are less of a concern. The reason that RBER is still a widely used metric for flash reliability is that it can be measured easily for raw flash chips and then be used as an indicator for the likelihood of experiencing UEs when using these chips inside an SSD. Projecting the rate of uncorrectable errors based on RBER was first proposed by Mielke *et al.* [32]. Since then it has become a commonly used technique [36]–[40] that is used, for example, to determine the appropriate strength of error correction in order to keep the rate of uncorrectable errors below a certain threshold.

This section studies the relationship between RBER and UEs as well as other types of errors, based on the Google data [4]. This question is motivated by the fact that individual cell errors that accumulate and become uncorrectable are only one possible cause of UEs. Other possible causes are defects, such as interconnect shorts that are created when applied voltage over time breaks down a marginal circuit, or bugs in the controller or firmware. The question was also motivated by Fig. 5, which to our surprise does not show a significant correlation between uncorrectable errors and RBER (recall the group of bars on the right in the figure).

We study the relationship between RBER and UEs in more detail in Fig. 15(a), which plots the median RBER and the fraction of drive days with UEs for all first generation disk models.<sup>1</sup> We see no correlation between RBER and

<sup>1</sup>Some of the 16 models in the figure were not included in Table 4, as they do not have enough data for some other analyses in the paper. The red markers correspond to SLC drives, the black ones to MLC drives.

UEs. We repeated the same analysis for the 95th percentile, rather than RBER, but did not see any evidence of a correlation either. Recall from Section II-D that all first generation drive models employ identical ECC, so if drives with different RBER still experience similar incidence of UEs it is not due to differences in the ECC.

To further study this issue we looked at the data at different granularities. Fig. 15(b) replots the same information as Fig. 15(a), but plots a separate data point for each individual drive. (The sample plot in the figure is for drives of model MLC-C.) Again, we see no indication that drives with higher RBER are more likely to experience uncorrectable errors.

Finally, we perform an analysis at an even finer time granularity, and study whether drive months with higher RBER are more likely to be months that experience a UE. Again we find no correlation.

We also studied the relationship between RBER and a number of other types of errors that Google's drives report (e.g., errors because an operation timed out, errors in accessing drive metadata, operations that succeeded only after a few retries), in particular whether RBER is higher in a month that also experiences other types of errors. We find that correlation coefficients are even lower for other error types.

In summary, we conclude that per-drive RBER is a poor predictor of UEs or other types of errors seen in the field. This might imply that the failure mechanisms leading to UEs are typically not due to individual cell errors (which RBER captures), but rather other mechanisms, such as flash defects or bugs in the controller or firmware.

## IX. LIMITATIONS AND FUTURE WORK

The three recent field studies described in this work provide a first step toward a better understanding of the reliability characteristics of NAND-based SSDs in the field. However, they also have some limitations and leave a number of interesting avenues open for future work.

Any future work that confirms trends and observations based on new field data sets would be valuable. Also, repeating prior analyses with an emphasis on controlling for confounding factors would be useful. For example, the Google analysis strives to control for confounding factors by confirming that trends persist when only looking at the data of individual clusters, or making sure that for a given model only drives that were put into production within the same year were considered. However, a year's difference might already be large for NAND and one might want to try and obtain data with manufacturing dates of drives and repeat the analysis with smaller manufacture date windows. This would rule out vintage effects, like those previously observed for HDDs [41], to be a major factor.

The existing work provides only little insight into the root cause of UEs. The observation in the Google study that there is little correlation between RBER and UEs points

towards defects or controller/firmware as contributors to UEs. It would be interesting to study patterns of UEs in more detail to gather additional evidence. For example, one could look for program status failures, which are often caused by defects, or for correlations to cluster or workload type or simultaneous batches of UEs across multiple sectors or pages. These could all be used to further rule out RBER-type causes.

Also, the results published so far make it difficult to distinguish NAND issues and SSD design issues. For example, the papers includes little detail about drive internal protection mechanisms, including for example the strength of the ECC, or whether drives include redundant NAND dies for RAID-like recovery of post-ECC errors. Also, more details on repair actions or results from failure analyses that might be performed by the manufacturer on returned devices could provide additional insights. For example, if reformatting a drive as part of a repair process resolved a problem, the problem was unlikely due to NAND issues.

Finally, it would be interesting to see a more detailed analysis of bad blocks and bad chips. For example, the Google study observes a long tail in the number of bad blocks per drive. It would be insightful to study the drives in the long tail more closely. For example, are all the bad blocks in the same bad chip that is causing the problem, or are they more scattered pointing to maybe workload factors or controller issues? Similarly for bad chips, did they go bad because lots of its NAND cells are wearing out or because of defects?

## X. SUMMARY

This paper provides a number of interesting insights into flash reliability in production use, based on three recent field studies [3]–[5]. Some of these support common assumptions and expectations, while many were unexpected. The summary below focuses on the more surprising results and implications.

- RBER, the standard metric for drive reliability, is not a good predictor of those failure modes that are the major concern in practice. In particular, in the field higher per-drive RBER does not translate to a higher incidence of uncorrectable errors. This might indicate that a common root cause of UEs in the field are defects or firmware/controller bugs, rather than single cell errors that accumulate.
- UBER, the standard metric to measure uncorrectable errors, is not very meaningful for field measurements. We see no correlation between UEs and number of reads, so normalizing uncorrectable errors by the number of bits read will artificially inflate the reported error rate for drives with low read count.
- Uncorrectable errors are a significant concern in practice: Depending on the model 20% to more than 90% of drives experience uncorrectable errors and two to six out of 1000 device days are affected.



- Both RBER and the number of uncorrectable errors grow with P/E cycles. While the literature contains reports of sublinear, linear, and superlinear growth of RBER with P/E cycles, we observe linear growth. We also observe no sudden spikes once a drive exceeds the vendor's P/E cycle limit, within the P/E cycle ranges we observe in the field.
- SLC drives, which are targeted at the enterprise market and considered to be higher end, are not more reliable than the lower end MLC drives with respect to uncorrectable errors for the P/E cycle ranges within the MLC cycle limits.
- The effect of temperature is more complex than one might expect. Rather than continuously increasing rates of uncorrectable errors with increasing temperature, some drives show stable or decreasing error rates under higher temperature. Reasons might be drive internal protection mechanisms that throttle drive operation under higher temperatures.
- Some drive models exhibit infant mortality that can be divided into two phases: A first phase with increasing error rates as bad cells are detected and removed from usage and a second phase with decreasing failure rates as weak cells have been weeded out.
- We observe that chips with smaller feature size tend to experience higher RBER, but are not necessarily the ones with the highest incidence of nontransparent errors.
- While flash drives offer lower field replacement rates than HDDs, they have a higher rate of problems that can impact the user, such as uncorrectable errors.
- Field replacement rates are often higher than what vendor specifications of MTTF might indicate.
- Previous errors of various types are predictive of later uncorrectable errors.
- Drives tend to either have less than a handful of bad blocks, or a large number of them, suggesting that impending chip failure could be predicted based on prior number of bad blocks (and maybe other factors). Also, a drive with a large number of factory bad blocks has a higher chance of developing more bad blocks in the field, as well as certain types of errors.

## APPENDIX

**Table 3** Summary of the Key Characteristics of the SSDs in the Facebook Study [3]

Platform	SSDs	PCIe	Per SSD				
			Capacity	Age (years)	Data written	Data read	UBER
A	1	v1, ×4	720 GB	2.4 ± 1.0	27.2 TB	23.8 TB	5.2 × 10 <sup>-10</sup>
B	2				48.5 TB	45.1 TB	2.6 × 10 <sup>-9</sup>
C	1	v2, ×4	1.2 TB	1.6 ± 0.9	37.8 TB	43.4 TB	1.5 × 10 <sup>-10</sup>
D	2				18.9 TB	30.6 TB	5.7 × 10 <sup>-11</sup>
E	1		3.2 TB	0.5 ± 0.5	23.9 TB	51.1 TB	5.1 × 10 <sup>-11</sup>
F	2				14.8 TB	18.2 TB	1.8 × 10 <sup>-10</sup>

**Table 4** Overview of Drive Models in the Google Study [4]

Model name	MLC-A	MLC-B	MLC-C	MLC-D	SLC-A	SLC-B	SLC-C	SLC-D	eMLC-A	eMLC-B
Generation	I	I	I	I	I	I	I	I	2	2
Vendor	I	II	I	I	I	I	III	I	I	IV
Flash type	MLC	MLC	MLC	MLC	SLC	SLC	SLC	SLC	eMLC	eMLC
Lithography (nm)	50	43	50	50	34	50	50	34	25	32
Capacity	480GB	480GB	480GB	480GB	480GB	480GB	480GB	960GB	2TB	2TB
PE cycle limit	3,000	3,000	3,000	3,000	100,000	100,000	100,000	100,000	10,000	10,000
Avg. PE cycles	730	949	529	544	860	504	457	185	607	377

**Table 5** Summary of the Key Characteristics of the SSDs in the Microsoft Study [5].  $\mu_{age}$ ,  $\mu_{reads}$ ,  $\mu_{writes}$  Refers to the Average Age, and Amount of Data Read/Written on Average Per Disk

Model	Size	$\mu_{age}$	$\mu_{reads}$	$\mu_{writes}$	Lith.
1-A	160GB	3.17 yrs.	NA	42.8 TB	34nm
1-B	160GB	3.31 yrs.	138.7 TB	25.1 TB	25nm
1-C	160GB	2.69 yrs.	99.9 TB	11.7 TB	25nm
1-D	480GB	1.92 yrs.	145.2 TB	40.3 TB	25nm
2-A	480GB	1.8 yrs.	NA	NA	20nm

## Acknowledgments

The authors would like to thank the reviewers, in particular, N. Mielke, who provided incredibly detailed feedback on the paper and many of their questions that arose while making edits for the final version of the paper. They would also like to thank the Platforms Team at Google, as well as N. Janevski and W. Chen for help with the data collection, and C. Sabol, T. Jeznach, and L. Barroso for feedback on earlier drafts of the paper. Finally, the first author would like to thank the Storage Analytics team at Google for hosting her in summer 2015 and for all their support. ■

## REFERENCES

- [1] JEDEC. (2016). *Solid State Drive (SSD) Requirements and Endurance Test Method*. [Online]. Available: <https://www.jedec.org/standards-documents/results/jesd218>
- [2] JEDEC. (2012). *Solid State Drive (SSD) Endurance Workloads*. [Online]. Available: <https://www.jedec.org/standards-documents/results/jesd219>
- [3] J. Meza, Q. Wu, S. Kumar, and O. Mutlu, "A large-scale study of flash memory failures in the field," in *Proc. ACM SIGMETRICS Int. Conf. Meas. Modelling Comput. Syst. (SIGMETRICS)*, 2015, pp. 177–190.
- [4] B. Schroeder, R. Lagisetty, and A. Merchant, "Flash reliability in production: The expected and the unexpected," in *Proc. 14th USENIX Conf. File Storage Technol. (FAST)*, Feb. 2016, pp. 67–80.
- [5] I. Narayanan et al., "SSD failures in datacenters: What? When? And why?" in *Proc. 9th ACM Int. Syst. Storage Conf. (SYSTOR)*, 2016, pp. 7:1–7:11.
- [6] N. Mielke, R. Frieki, I. Kalistirsky, M. Quan, D. Ustinov, and V. Vasudevan, "Reliability of solid-state drives based on NAND flash memory," *Proc. IEEE*, 2017, DOI: 10.1109/JPROC.2017.2725738.
- [7] R. Micheloni, L. Crippa, and A. Marelli, *Inside NAND Flash Memories*. Berlin, Germany: Springer-Verlag, 2010.
- [8] L. N. Bairavasundaram, G. R. Goodson, B. Schroeder, A. C. Arpaci-Dusseau, and R. H. Arpaci-Dusseau, "An analysis of data corruption in the storage stack," in *Proc. 6th USENIX Conf. File Storage Technol. (FAST)*, San Jose, CA, USA, Feb. 2008, p. 8.
- [9] L. Mearian. (2009). *Intel Confirms Data Corruption Bug in New SSDs, Halts Shipments*. [Online]. Available: <http://www.computerworld.com/article/2526707/data-center/intel-confirms-data-corruption-bug-in-new-ssds-halts-shipments.html>
- [10] B. Ferreira. (2015). *Some Samsung SSDs May Suffer From a Buggy TRIM Implementation*. [Online]. Available: <http://techreport.com/news/28473/some-samsung-ssds-may-suffer-from-a-buggy-trim-implementation>
- [11] M. Campbell. (2015). *Apple Fixes 2015 MacBook Pro Flash Storage Issue in Firmware Update*. [Online]. Available: <http://appleinsider.com/articles/15/07/22/apple-fixes-2015-macbook-pro-flash-storage-issue-in-firmware-update>
- [12] A. Chanthadavong. (2016). *Amazon Web Services Sydney Suffers Outage*. [Online]. Available: <http://www.zdnet.com/article/amazon-web-services-sydney-suffers-outage>
- [13] Z. Whittaker. (2012). *Amazon Explains Latest Cloud Outage: Blame the Power*. [Online]. Available: <http://www.zdnet.com/article/amazon-explains-latest-cloud-outage-blame-the-power/>
- [14] M. Zheng, J. Tucek, F. Qin, and M. Lillibridge, "Understanding the robustness of SSDs under power fault," in *Proc. 11th USENIX Conf. File Storage Technol. (FAST)*. San Jose, CA, USA, 2013, pp. 271–284.
- [15] H.-W. Tseng, L. M. Grupp, and S. Swanson, "Understanding the impact of power loss on flash memory," in *Proc. 48th Design Autom. Conf.*, 2011, pp. 35–40.
- [16] C. Zambelli, P. King, P. Olivo, L. Crippa, and R. Micheloni, "Power-supply impact on the reliability of mid-1X TLC NAND flash memories," in *Proc. IEEE Int. Rel. Phys. Symp. (IRPS)*, Apr. 2016, pp. 2B-3-1–2B-3-6.
- [17] A. Klein. (2016). *One Billion Drive Hours and Counting: Q1 2016 Hard Drive Stats*. [Online]. Available: <https://www.backblaze.com/blog/hard-drive-reliability-stats-q1-2016/>
- [18] J. Kim, E. Lee, J. Choi, D. Lee, and S. H. Noh, "Chip-Level RAID with flexible stripe size and parity placement for enhanced SSD reliability," *IEEE Trans. Comput.*, vol. 65, no. 4, pp. 1116–1130, Apr. 2016.
- [19] Intel. (2011). *Validating the Reliability of Intel Solid-State Drives*. [Online]. Available: <http://www.intel.de/content/dam/doc/technology-brief/intel-it-validating-reliability-of-intel-solid-state-drives-brief.pdf>
- [20] B. Schroeder and G. A. Gibson, "Disk failures in the real world: What does an MTTF of 1,000,000 hours mean to you?" in *Proc. FAST*, vol. 7. 2007, pp. 1–16.
- [21] E. Pinheiro, W.-D. Weber, and L. A. Barroso, "Failure trends in a large disk drive population," in *Proc. FAST*, vol. 7. 2007, pp. 17–23.
- [22] L. N. Bairavasundaram, G. R. Goodson, S. Pasupathy, and J. Schindler, "An analysis of latent sector errors in disk drives," in *Proc. ACM SIGMETRICS Int. Conf. Meas. Modelling Comput. Syst.*, New York, NY, USA, 2007, pp. 289–300.
- [23] D. Wei, L. Deng, P. Zhang, L. Qiao, and X. Peng, "A page-granularity wear-leveling (PGWL) strategy for NAND flash memory-based sink nodes in wireless sensor networks," *J. Netw. Comput. Appl.*, vol. 63, pp. 125–139, Mar. 2016.
- [24] Y. Cai, Y. Luo, E. F. Haratsch, K. Mai, and O. Mutlu, "Data retention in MLC NAND flash memory: Characterization, optimization, and recovery," in *Proc. 21st IEEE Int. Symp. High Perform. Comput. Archit. (HPCA)*, Burlingame, CA, USA, Feb. 2015, pp. 551–563.
- [25] Y. Cai, O. Mutlu, E. F. Haratsch, and K. Mai, "Program interference in MLC NAND flash memory: Characterization, modeling, and mitigation," in *Proc. IEEE 31st Int. Conf. Comput. Design (ICCD)*, Oct. 2013, pp. 123–130.
- [26] X. Xu and H. H. Huang, "Exploring data-level error tolerance in high-performance solid-state drives," *IEEE Trans. Reliability*, vol. 64, no. 1, pp. 15–30, Mar. 2015.
- [27] M. Huang, Z. Liu, and L. Qiao, "Asymmetric programming: A highly reliable metadata allocation strategy for MLC NAND flash memory-based sensor systems," *MDPI Sensors*, vol. 14, no. 10, pp. 18851–18877, 2014.
- [28] X. Jimenez, D. Novo, and P. Ienne, "Wear unleveling: Improving NAND flash lifetime by balancing page endurance," in *Proc. 12th USENIX Conf. File Storage Technol. (FAST)*, Santa Clara, CA, USA, 2014, pp. 47–59.
- [29] L. Zuolo, C. Zambelli, P. Olivo, and A. Marelli, "LDPC soft decoding with reduced power and latency in 1X-2X NAND flash-based solid state drives," in *Proc. Int. Memory Workshop (IMW)*, May 2015, pp. 1–4.
- [30] L. M. Grupp et al., "Characterizing flash memory: Anomalies, observations, and applications," in *Proc. 42nd Annu. IEEE/ACM Int. Symp. Microarchit.*, New York, NY, USA, Dec. 2009, pp. 24–33.
- [31] H. Sun, P. Grayson, and B. Wood, "Quantifying reliability of solid-state storage from multiple aspects," in *Proc. SNAPI*, 2011.
- [32] N. Mielke et al., "Bit error rate in NAND flash memories," in *Proc. IEEE Int. Rel. Phys. Symp.*, Apr./May 2008, pp. 9–19.
- [33] L. M. Grupp, J. D. Davis, and S. Swanson, "The bleak future of NAND flash memory," in *Proc. 10th USENIX Conf. File Storage Technol.*, Berkeley, CA, USA, 2012, p. 2.
- [34] N. Mielke, H. P. Belgal, A. Fazio, Q. Meng, and N. Righos, "Recovery effects in the distributed cycling of flash memories," in *Proc. IEEE 44th Annu. Int. Rel. Phys. Symp.*, Mar. 2006, pp. 29–35.
- [35] N. Papandreou et al., "Effect of read disturb on incomplete blocks in MLC NAND flash arrays," in *Proc. IEEE 8th Int. Memory Workshop (IMW)*, May 2016, pp. 1–4.
- [36] Y. Cai et al., "Flash correct-and-refresh: Retention-aware error management for increased flash memory lifetime," in *Proc. IEEE 30th Int. Conf. Comput. Design (ICCD)*, Sep./Oct. 2012, pp. 94–101.
- [37] R.-S. Liu, C.-L. Yang, and W. Wu, "Optimizing NAND flash-based SSDs via retention relaxation," in *Proc. 10th USENIX Conf. File Storage Technol. (FAST)*, San Jose, CA, USA, Feb. 2012, p. 11.
- [38] M. Balakrishnan, A. Kadav, V. Prabhakaran, and D. Malkhi, "Differential RAID: Rethinking RAID for SSD reliability," *Trans. Storage*, vol. 6, no. 2, pp. 4:1–4:22, Jul. 2010.
- [39] C. Zambelli et al., "A cross-layer approach for new reliability-performance trade-offs in MLC NAND flash memories," in *Proc. Design, Autom. Test Eur. Conf. Exhib.*, San Jose, CA, USA, Mar. 2012, pp. 881–886.
- [40] G. Wu, X. He, N. Xie, and T. Zhang, "Exploiting workload dynamics to improve ssd read latency via differentiated error correction codes," *ACM Trans. Des. Autom. Electron. Syst.*, vol. 18, no. 4, pp. 55:1–55:22, Oct. 2013.
- [41] J. G. Elerath, "Specifying reliability in the disk drive industry: No more MTBF's," in *Proc. Annu. Rel. Maintainability Symp.*, Jan. 2000, pp. 194–199.

## ABOUT THE AUTHORS

**Bianca Schroeder** received the Ph.D. degree from the Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, USA, under the direction of M. Harchol-Balder.

She is an Associate Professor and Canada Research Chair in the Computer Science Department, University of Toronto, Toronto, ON, Canada, also currently serving as an Associate Department Chair at the Computer and Mathematical Sciences Department, University of Toronto, Scarborough. Before



joining UofT, she spent two years as a Postdoctoral Researcher at Carnegie Mellon University, working with G. Gibson.

Prof. Schroeder is an Alfred P. Sloan Research Fellow, the recipient of the Outstanding Young Canadian Computer Science Prize of the Canadian Association for Computer Science, an Ontario Early Researcher Award, an NSERC Accelerator Award, a two-time winner of the IBM PhD fellowship and her work has won four best paper awards and one best presentation award. She has served on numerous conference program committees and has cochaired the TPCs of Usenix FAST-14, ACM Sigmetrics-14, and IEEE NAS-11.

**Arif Merchant** received the B.Tech. degree from the Indian Institute of Technology Bombay (IIT Bombay), Powai, Mumbai, India and the Ph.D. degree in computer science from Stanford University, Stanford, CA, USA.

He is a Research Scientist at Google Inc., Mountain View, CA, USA, and leads the Storage Analytics group, which studies interactions between components of the storage stack. His interests include distributed storage systems, storage management, and stochastic modeling.

Dr. Merchant is an Association for Computing Machinery (ACM) Distinguished Scientist.



**Raghav Lagisetty** received the B.S. degree in computer science from the Indian Institute of Technology Bombay (IIT Bombay), Powai, Mumbai, India and the M.S. degree in computer science from the University of Arizona, Tucson, AZ, USA.

He is an Engineering Lead at Google Inc., Mountain View, CA, USA, for the backend infrastructure team of Ads systems. Previously he has built ground up products and engineering teams in the areas of cloud storage virtualization, storage reliability, and big data analytics.

