Machine Learning Models for SSD and HDD Reliability Prediction

Riccardo Pinciroli, Ph. D., Gran Sasso Science Institute

Lishan Yang, William & Mary

Jacob Alter, Ipsos North America

Evgenia Smirni, Ph. D., William & Mary

Key Words: Supervised learning, Classification, Data centers, Storage devices, SSD, HDD

2022 Annual Reliability and Maintainability Symposium (RAMS) | 978-1-6654-2432-5/22/$31.00 ©2022 IEEE | DOI: 10.1109/RAMS51457.2022.9893942

SUMMARY & CONCLUSIONS

This paper compares hard disk drives (HDDs) and solid-state drives (SSDs), the two most used storage devices in data centers, which frequently fail and are among the main causes of data center downtime. Using a six-year field data of 100,000 HDDs from the Backblaze dataset and a six-year field data of 30,000 SSDs from a Google data center, we characterize workload conditions that prompt drive failures. %and show that they differ from common expectation. We develop machine learning models that accurately predict the drive failure state several days in advance and provide highly interpretable results that are useful to identify the causes and symptoms of drive failures.

1 INTRODUCTION

Data centers dependability is highly affected by storage devices such as hard disk drives (HDDs) and solid state drives (SSDs) [1,2]. Accurate prediction of storage device failures increases the data center dependability by enabling actions to reduce data loss [3], such as drive replacement before failures occur, migration of data to other resources, and allocation of virtual machines to disks that are not prone to failure [4]. Observations from HDDs (SSDs) cannot be generalized to SSDs (HDDs) due to their different physical mechanisms. Previous research on HDD dependability analyzes data collected during a period of up to two years from a few disk models [5,6]. Most SSD studies use simulated or controlled environments [7,8] and consider specific error types [9,10].

In this paper, we investigate HDD and SSD failures by analyzing disk logs collected from real-world data centers over a period of six years. We analyze Self-Monitoring, Analysis and Reporting Technology (SMART) traces of HDDs from the Backblaze data center [11] and the logs of SSDs collected at a Google data center. We study the role of errors accounted in the logs in triggering future drive failures.

The datasets considered in this paper contain a tremendous amount of data (i.e., tens of millions of daily drive reports) regarding drive performance and errors. Error logs are analyzed to find attributes that are related to drive malfunctions and may be used to predict forthcoming failures. Statistical methods do not achieve good results and there is no evidence that the repairing process is started by deterministic decision rules. We train machine learning models with monitoring logs to achieve accurate and fast failure prediction several days before the drive failure.

Highly imbalanced datasets (the healthy-faulty ratio is 10,000:1) make it difficult to achieve simultaneously high true positive rates and low false positive ones. We present different ways to partition the HDD and SSD datasets to increase model accuracy. This partitioning is based on workload analysis that was first developed in [12] for SSDs and focuses on the discovery of certain drive attributes. We saw that similar partitioning can be also successfully applied for the case of HDDs. We also focus on the interpretability of the machine learning models and derive insights that can be used to drive proactive disk management policies. Our findings are summarized as follows:

- Drive failures are triggered by a set of attributes. There is no single metric that triggers a drive failure after it reaches a certain threshold.
- Several machine learning predictors are quite successful for failure prediction. Random Forest models are found to be the most successful for both SSDs and HDDs.
- Datasets may be partitioned to improve the performance of the classifier.

Partitioning SSDs on the drive age attribute and HDDs on head flying hours (i.e., SMART 240) increases model accuracy.

2 SSD AND HDD DATASET

The SSD dataset consists of daily performance logs for three multi-level cell (MLC) SSD models collected at a Google data center over a period of six years. All three models are manufactured by the same vendor and have a 480GB capacity and a lithography on the order of 50nm. They utilize custom firmware and drivers, meaning that error reporting is done in a proprietary format rather than through standard SMART features. We refer to the three models as MLC-A, MLC-B, and MLC-D in accordance with the naming in [9,13]. We have data on over 10,000 unique drives for each drive model, totaling over 40,000,000 daily drive reports overall. The logs used in this paper report daily summaries of drive activity. Drives are uniquely identified by their drive ID, which is a hashed value of their serial number. For each day of operation, the following

Authorized licensed use limited to: Gran Sasso Science Institute. Downloaded on April 18,2023 at 10:13:49 UTC from IEEE Xplore. Restrictions apply.

metrics are reported: drive age; number of operations and program-erase cycles; drive failure state; number of bad blocks in the drive; count of different errors (e.g., correctable and uncorrectable errors).

The HDD dataset contains daily logs of HDDs collected at a Backblaze data center over six years. Basic information such as serial number, model, and date is reported, as well as standard SMART features. We consider the 5 most popular Seagate HDD models starting from January 17, 2014, until December 31, 2019. There are over 100,000 unique HDDs in the dataset, with more than 100,000,000 daily reports in total. HDDs are uniquely identified by their serial number. A snapshot operation is performed every day for all operational hard drives.

3 DRIVE FAILURE AND CAUSES

In this section, we analyze failures of SSDs and HDDs. We identify possible symptoms and causes of those failures.

3.1 *Drive Failure*

Table 1 shows the percentage of failures for each drive model in SSD and HDD datasets. On average, the failure rate of SSDs is higher than the one of HDDs. Table 2 provides more insights by reporting statistics on the frequency of failures for the same drive. Unexpectedly, we find that some SSDs have failed as many as four times over the course of their lifetime. Nonetheless, 89.6% of drives with failures fail only once. HDDs have no more than 2 failures per drive, and only 0.194% of the failed drives has 2 failures.

Table 1 – High-level failure incidence statistics.

	Model	#Failures	%Failed
SSD	MLC-A	734	6.95
	MLC-B	1565	14.3
	MLC-D	1580	12.5
	All	3879	11.29
HDD	ST12000NM0007	1448	3.76
	ST3000DM001	1357	31.89
	ST4000DM000	3710	10.21
	ST8000DM002	354	3.47
	ST8000NM0055	435	2.91
	All	7743	7.01

Table 2 – Distribution of lifetime failure counts.

	# of Failures	% of Drives	% of Failed Drives
SSD	0	88.71	—
	1	10.10	89.60
	2	1.038	9.208
	3	0.133	1.180
	4	0.001	0.001
HDD	0	93.35	—
	1	6.64	99.81
	2	0.01	0.194

HDD failure events are directly logged using the failure feature, while failures are not directly reported for SSDs.

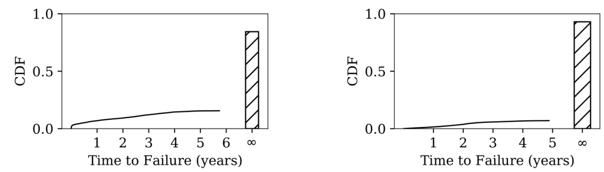

(a) SSD (b) HDD

Figure 1 - CDF of the length of the drive's operational period. The bar indicates what proportion of operational periods are not observed to end. Failure rate of SSDs is larger than HDDs

Hence, in the SSD case, we define failure events as a drive's last day of operational activity This is a natural point of failure since, after this point in the timeline, the drive has ceased normal function and needs to be repaired. If no performance summaries are documented in the SSD log, or substantially higher rates of inactivity relatively to normal drive operation are discovered, i.e., an absence of read or write operations provisioned to the drive, then we define a failure as happening directly prior to this period of inactivity, if such a period exists. For both HDDs and SSDs, the failed drive may or may not be repaired successfully. If it is repaired, it re-enters its operational period.

Fig. 1 presents the CDF of the length of operational periods (alternately denoted "time to failure"). The CDF includes both operational periods starting from the beginning of the drive's lifetime and operational periods following a post-failure reentry if the failed drive is successfully repaired. More than 80% and 90% of operational periods of SSDs and HDDs, respectively, are not observed to end with failure during the 6 year sampling period; this probability mass is indicated by the bar centered at infinity. The figure indicates that there is a substantial variability in the drive operational time, with the majority of operating times being long. Yet, there is a nonnegligible portion of operating times that are interrupted by failures.

Comparing SSDs with HDDs in Figs. 1(a) and (b), the maximum time to failure of HDDs is less than 5 years, while the one of SSDs is close to 6 years. Only 7% of HDDs fail after an operational period, while for SSDs it is near to 20%. This is due to 1) the lower failure rate of HDDs compared to the one of SSDs and 2) the different maximum number of failures on a single drive (i.e., 2 for HDDs and 4 for SSDs, see Table 2).

3.2 *Symptoms and Causes of Failures*

Since the features/information provided by the two datasets are different, here we analyze SSDs and HDDs separately.

1) Age in SSDs: Recall that Table 1 shows that, among the drives represented in the datasets, 11.29% of SSDs and 7.01% of HDDs fail at least once. A natural question is when do these failures (and the following swaps for SSDs) occur in the drive's lifetime: what is the role of age in drive failure? Figure 2(a) reports the CDF of the failure age (solid line) as a function of the drive age for SSDs. The figure shows that there are many more drive failures in the first 90 days of drive operation than at any other point in the drive lifetime. In fact, 15% of observed failures occur on drives less than 30 days old and 25% occur on drives less than 90 days old. This seems to

Authorized licensed use limited to: Gran Sasso Science Institute. Downloaded on April 18,2023 at 10:13:49 UTC from IEEE Xplore. Restrictions apply.

indicate that these drives have an infancy period during which drive mortality rate is particularly high. This performance pattern has been noticed previously in similar studies of SSDs in the wild [14].

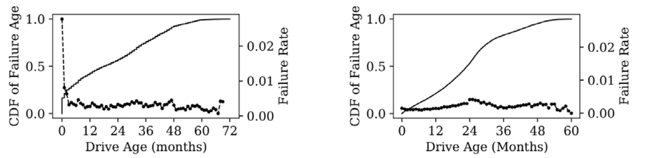

(a) SSD (b) HDD

Figure 2 - The CDF of the age of failed drives (solid line) and the proportion of functioning drives that fail at a given age level, in months (dashed line)

The slope of the CDF in Figure 2(a) gives us an estimate of the rate at which swaps occur at a given drive age. However, this estimate is skewed since not all drive ages are equally represented in the data. For example, the rate of failures seems to slow down following the four year mark, but this is due to the fact that drives of this age level are not as common in the data. We normalize the number of swaps within a month by the amount of drives represented in the data at that month to produce an unbiased failure rate for each month (dashed line in Figure 2). We see that this rate evens out after the third month, indicating that the length of this observed highfailure infancy period is approximately 90 days. Accordingly, for the remainder of this paper, we distinguish drive swaps as young versus old, i.e., those swaps occurring before vs. after the 90-day mark. Beyond the 90-day mark, we observe that the failure rate is roughly constant, suggesting that, even if drives become very old, they are not more prone to failure.

One potential explanation for the spike in failures for infant drives is that they are undergoing a “burn-in” period in manufacturing malfunctions. This is a common practice in data centers, wherein new drives are subjected to a series of high-intensity workloads in order to test their resilience and check for manufacturing faults. These increased workloads could stress the drive, leading to a heightened rate of failure.

2) Head Flying Hours in HDDs: Similar to the drive age in SSDs shown in Figure 2(a), we also present a similar plot for HDDs in Figure 2(b). Differently from the high failure rate of young SSDs, the failure rate for HDDs regarding drive age is relatively small (less than 1%). Therefore, we need to find other features which may be related to failure rate.

We examine all SMART features for HDDs, and find out that head flying hours (HFH, SMART 240) is highly related to failures. Here we define two kinds of disks regarding head flying hours with a certain threshold: i) Large HFH disks are observed at least once with head flying hours larger than the threshold; ii) Small HFH disks always have head flying hours smaller than or equal to the threshold.

Figure 3(a) shows the failure rate of small and large HFH drives as a function of the threshold. The average faialure rate of all HDDs is also reported (baseline, see dashed line). We observe two situations with high failure rate in Figure 3: i) small HFH when threshold is less than 3000 (the beginning of small HFH line), and ii) large HFH when threshold is larger than 40,000 (the end of large HFH line).

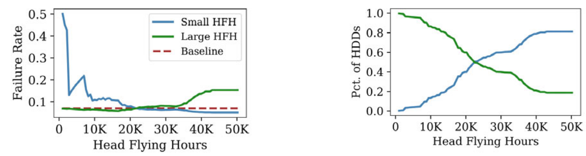

(a) Failure Rate (b) HDD Percentage

Figure 3 - Failure rate and percentage of corresponding HDDs with different head flying hours (SMART 240).

The percentage of HDDs if we partition the dataset according to HFH is also shown in Figure 3(b). When the threshold is less than 3000, the percentage of small HFH drives is less than 3%, therefore it is not representative. When the threshold is larger than 40,000, the percentage of large HFH is about 20%, which is worth to be considered. The failure rate of these 20% large HFH disks is 17%, which is much higher than the 7% failure average (baseline). Balancing the failure rate and percentage of large HFH disks, we use threshold = 40;000 to split the dataset. This observation (HDDs with small and large HFH have different resilience behavior) guides us to split the dataset for better prediction, see Section 4.3 for more details.

4 *FAILURE PREDICTION*

In this section, we use machine learning models to detect SSD and HDD failures that will occur within $N \geq 0$ days. Causes of SSD and HDD failure investigated in Section 3.2 are used to improve model accuracy. Feature importance allows identifying those attributes that are more critical for SSD and HDD lifetimes.

4.1 *Model description*

Input. Both SSD and HDD datasets present daily statistics and are extremely imbalanced, i.e., the number of healthy disks (majority class) is larger than the number of faulty ones (minority class). In the SSD case, the ratio of healthy and defective drives is 10,000:1, for HDDs it is 13,000:1. To deal with such imbalanced datasets, we under-sample the majority class, use cross-validation for training and testing the model, and evaluate its performance with measures that are not affected by imbalanced datasets.

Under-sampling. The training set of both datasets is undersampled to result in a 1:1 healthy-faulty drives ratio to avoid the classifier being biased toward healthy drives. We use a random strategy, i.e., observations to be removed are randomly chosen from the majority class. We observe that the model performance is not profoundly affected by considering different under-sampling strategies and healthy-faulty ratios.

Cross-Validation. Classifiers are cross-validated by splitting each dataset into five different folds (4 folds for training, 1 for testing). The dataset is partitioned by drive ID (i.e., all observations of a drive belong to the same fold and are not used concurrently for training and testing). If folds are

Authorized licensed use limited to: Gran Sasso Science Institute. Downloaded on April 18,2023 at 10:13:49 UTC from IEEE Xplore. Restrictions apply.

created by randomly partitioning the dataset as in [5,9], future observations of a drive may be used to predict the past failure state of the same drive. This is undesirable since no future information is available in real scenarios.

Output. The model returns a continuous value in the interval (0,1), i.e., the probability that the drive fails. Such a value is discretized by using a discrimination threshold, α. If the output is larger than α, then the model predicts a failure; otherwise, the model predicts a non-failure. Due to its insensitiveness to imbalanced datasets, receiver operating characteristic (ROC) is a widely used metric to evaluate the accuracy of binary classifiers [9]. ROC curve plots the true positive rate, TPR=TP/(TP+FN), against the false positive rate, FPR=FP/(FP+TN), of the analyzed classifier by considering different values of α. We also use the area under the ROC curve (i.e., AUROC) to determine the performance of the proposed classifier. The AUROC value is between 0.5 (i.e., random classifier) and 1 (i.e., perfect predictor).

4.2 *Prediction Accuracy*

We investigate and report the performance of different classification models in Table 3. It shows the AUROC of predictors for different lookahead windows and datasets (i.e., SSDs and HDDs). Although Table 3 shows only one value for each classifier, we investigate their performance with various hyperparameters. XGBoost and Random Forest models have similar performance, but the time required to train the latter classifier is only 5% of the time required for training the former one. Table 3 also shows that, for all classifiers, the AUROC decreases when the lookahead window increases.

Table 3 - AUROC for different predictors and lookahead windows, N. The cross-validated average AUROC is provided with the standard deviation across folds.

	N	0	1	7
SSD	Logistic Reg.	0:796±0:010	0:765±0:009	0:713±0:010
	k-NN	0:816±0:013	0:791±0:009	0:716±0:008
	SVM	0:821±0:014	0:795±0:011	0:728±0:011
	Neural network	0:857±0:007	0:828±0:004	0:770±0:008
	Decision tree	0:872±0:007	0:840±0:007	0:780±0:006
	XGBoost	0:904±0:001	0:873±0:002	0:809±0:001
	Random forest	0:905±0:008	0:859±0:007	0:803±0:008
HDD	Logistic Reg.	0:668±0:001	0:669±0:001	0:668±0:001
	k-NN	0:699±0:022	0:699±0:026	0:691±0:029
	SVM	0:679±0:012	0:689±0:009	0:684±0:009
	Neural network	0:684±0:065	0:683±0:076	0:682±0:076
	Decision tree	0:886±0:051	0:870±0:053	0:837±0:052
	XGBoost	0:904±0:001	0:888±0:001	0:854±0:001
	Random forest	0:903±0:013	0:888±0:010	0:854±0:006

Figure 4 plots the AUROC of the Random Forest prediction on HDD (solid line) and SSD (dashed line) datasets against different lookahead windows. Each value, obtained by averaging the AUROC of different cross-validation folds, is plotted with its standard deviation. In both cases, the Random Forest performance decreases for longer lookahead windows and better AUROC values are observed for the HDD dataset. Figure 4 suggests that the Random Forest can efficiently predict SSD and HDD failures for N ≤ 2 and N ≤ 8 days lookahead, respectively.

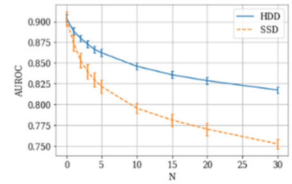

Figure 4 - Random Forest AUROC as a function of lookahead window size, N. Error bars indicate the standard deviation of the cross-validated error across folds.

4.3 *Model Improvement*

Section 3.2 shows that many SSD failures are related to the drive age, while HDD ones are affected by the head flying hours (i.e., SMART 240). Here, we use those attributes to improve the performance of the model. Each dataset is split based on the value of the considered feature (i.e., drive age or head flying hours). Then, a model is trained and validated on each sub-dataset and the performance of each new model is compared to the performance obtained without splitting the dataset.

Fig. 5 shows the ROC curve of the model when it is trained on different sub-dataset and predicts the state of each drive in the next 24 hours. As depicted in Fig. 5(a), the prediction model works better with young SSDs (i.e., drive age smaller than 3 months) since its AUROC is significantly larger (0.961) than the one shown in Fig. 4 (0.906). This comes at the expense of slightly reduced performance for older drives (0.894). Similarly, Fig. 5(b) shows the performance improvement observed by splitting the HDD dataset on the head flying hours feature (i.e., 40,000 hours). The model can better predict the state of HDDs that spend a longer time in positioning their heads (0.929). Improvements are observed comparing to the no-split strategy whose AUROC is 0.902, while the performance for drives with small head flying hours slightly decreases (0.890). It is worth noting that the 20% of swap-inducing failures in the SSD dataset are young failures, while the 25% of HDDs has a large head flying hours.

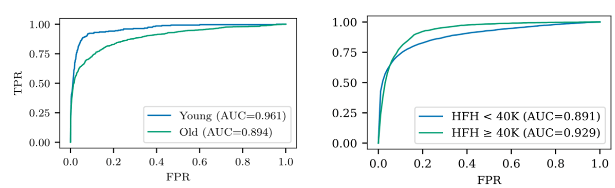

(a) SSD (b) HDD

Figure 5 - ROC curves after splitting the dataset on drive features. Prediction model is random forest with lookahead window N=0 days. HFH is for Head Flying Hours (i.e., SMART 240).

4.4 *Model Interpretability*

Random Forest models assign a score to each attribute based on its relevance for solving the classification problem. This increases the model interpretability since it is possible to identify those features that are more related to drive failures.

Authorized licensed use limited to: Gran Sasso Science Institute. Downloaded on April 18,2023 at 10:13:49 UTC from IEEE Xplore. Restrictions apply.

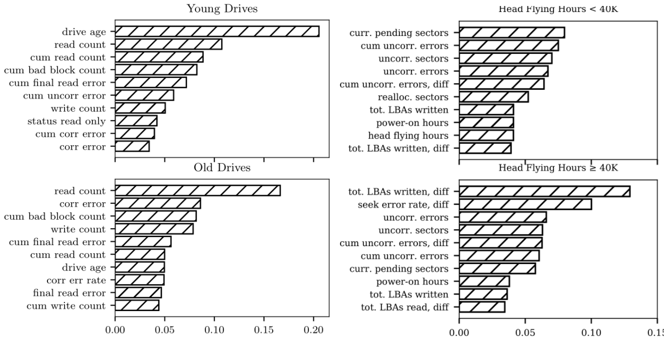

(a) Failure Rate (b) HDD Percentage

Figure 6 - Feature importance for Random Forest models after splitting the dataset on drive features.

Fig. 6 shows the TOP-10 features for each sub-dataset considered in Section 4.3 (i.e., young and old SSDs, HDDs with small and large head flying hours). Fig. 6(a) shows the feature ranking for the SSD dataset. When considering young drives, the drive age is the most important feature, followed by the read count, its cumulative value, and the cumulative number of bad blocks. For old SSDs, featu res counting correctable errors and read/write operations, and the cumulative number of bad blocks are in the TOP-4. It is expected that read and write counts are more relevant for the state prediction of old drives since young drives may have only a few activities at the failure time. Fig. 6(b) depicts the feature importance for the HDD dataset. The number of current pending sectors, uncorrectable errors, uncorrectable sectors, and r eallocated sectors are among the most important features for detecting failing drives. This is similar to what is observe d in [9]. The attribute ranking of HDDs with large head flying hours provides new insights. The two most relevant features are the incremental step of written logical block addressing (LBA) and seek error rate, followed by the number of uncorrectable errors and uncorrectable sectors. The seek error rate and the uncorrectable sector count are observed to be important features also in [15]. The reallocated sector count is not in the TOP-10 important features for HDDs with large head flying hours.

5 RELATED WORK

Prior work investigates the main components that affect the data center dependability [16,17,18,19] and storage drives are among the most important ones [1,2]. Different approaches are proposed to predict drives failures on various datasets [15,20,21]. However, to fairly compare the performance of our method with other approaches, the same dataset must be used [22]. The Backblaze dataset is used in [5,23] to train and validate their approaches for predicting disk failures while [9] uses the Backblaze and the Google datasets for investigating faulty HDDs and SSDs. The regularized greedy forest adopted by Botezatu et al. [5] has 98% precision and accuracy, but it splits observations from the same drive into training and test sets not considering error and workload correlations across different drive days. Similarly to Wang et al. [6], it also undersamples both training and test sets [23], but resampling the whole dataset (training and test sets) generates overoptimistic results. Aussel et al. [23] train and evaluate a random forest on a small subset of the Backblaze dataset (i.e., only data from 2014) resulting in high precision and considerable recall. However, they filter out observations with similar features and different failure states, that requires apriori knowledge of the drive state. Mahdisoltani et al. [9] investigates different prediction models to predict uncorrectable errors and bad blocks in HDDs and SSDs, and they show that random forests provide good performance for this task. Although our approach is similar to the one presented in [9], we aim to predict drive failures and do not try to predict errors. Approaches for online HDD failure prediction are also investigated [24,25] by using online random forests. The SSD dataset does not provide global timestamps, so online prediction cannot be implemented.

In this paper, we explore the capability of random forests for predicting drive failures and investigate possible enhancements by statistically analyzing drive features and using different models based on observed attributes. We consider a conceivably long lookahead window and use two large and real datasets for training and validating the proposed machine learning approach. To the best of our knowledge, storage device failures have never been studied by splitting the dataset based on attribute values.

6 CONCLUSION

In this paper, we investigate SSD and HDD failures using two traces from production environments. Daily logs for 30,000 SSDs are collected at a Google data center, while 100,000 HDDs are observed at a Backblaze data center. We train and test different classifiers to predict faulty SSDs and HDDs, and note that Random Forest models provide accurate predictions with a short training time. Their high interpretability makes them the best predictor for the analyzed problem. We observe that splitting each dataset based on attribute values of its observations allows increasing the performance of random forests. The drive age is a critical attribute for predicting SSD failures; drives failing before being three months old can be detected easier than other drives. When predicting faulty HDDs, a higer detection rate is observed for drives with head flying hours (SMART 240) longer than 40,000 hours.

ACKNOWLEDGEMENT

We deeply thank Arif Merchant of Google for making the SSD trace available to us. This material is based upon work supported by the National Science Foundation (NSF) grants (#1717532 and #CCF-1649087). Lishan Yang is partially supported by the CoVA CCI grant C-Q122-WM-02. Riccardo Pinciroli is supported by the MIUR PRIN project SEDUCE 2017TWRCNB.

REFERENCES

1. B. Schroeder, G. A. Gibson, "Disk failures in the real world: What does an MTTF of 1,000,000 hours mean to you?" *5th USENIX Conference on File and Storage Technologies,* 2007, pp. 1–16.
2. R. Birke et al., "(Big)data in a virtualized world: volume, velocity, and variety in cloud datacenters", *Proceedings of*

Authorized licensed use limited to: Gran Sasso Science Institute. Downloaded on April 18,2023 at 10:13:49 UTC from IEEE Xplore. Restrictions apply.

the 12th USENIX conference on File and Storage Technologies, 2014, pp. 177–189.
3. A. Ma et al., “RAIDShield: characterizing, monitoring, and proactively protecting against disk failures”, *ACM Transactions on Storage* 11.4 (2015), pp. 1–28.
4. J. Meza et al., “A Large-Scale Study of Flash Memory Failures in the Field”, *Proceedings of the 2015 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, 2015, pp. 177–190.
5. M. K. Qureshi et al., “AVATAR: A VariableRetention-Time (VRT) Aware Refresh for DRAM Systems”, *45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, 2015, pp. 427–437.
6. M. M. Botezatu et al., “Predicting disk replacement towards reliable data centers”, *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining,* 2016, pp. 39–48.
7. B. Schroeder et al., “Flash Reliability in Production: the Expected and the Unexpected”, *14th USENIX Conference on File and Storage Technologies,* 2016, pp. 67–80.
8. N. Aussel et al., “Predictive models of hard drive failures based on operational data”, *16th IEEE International Conference on Machine Learning and Applications,* 2017, pp. 619–625.
9. F. Mahdisoltani et al, “Proactive error prediction to improve storage system reliability”, *2017 USENIX Annual Technical Conference*, 2017, pp. 391–402.
10. B. Nie et al., “Characterizing Temperature, Power, and Soft-Error Behaviors in Data Center Systems: Insights, Challenges, and Opportunities”, *MASCOTS 2017, 2017,* pp. 22–31.
11. L. P. Queiroz et al., “A fault detection method for hard disk drives based on mixture of Gaussians and nonparametric statistics”, *IEEE Transactions on Industrial Informatics* 13.2 (2017), pp. 542–550.
12. B. Schroeder et al., “Reliability of NAND-based SSDs: What field studies tell us”, *Proceedings of the IEEE* 105.9 (2017), pp. 1751–1769.
13. G. Wang et al., “What Can We Learn from Four Years of Data Center Hardware Failures?”, *47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, 2017, pp. 25–36.
14. Y. Luo et al., “Improving 3D NAND Flash Memory Lifetime by Tolerating Early Retention Loss and Process Variation”. *Abstracts of the 2018 ACM International Conference on Measurement and Modeling of Computer Systems, SIGMETRICS 2018,* 2018, p. 106.
15. B. Nie et al., “Machine Learning Models for GPU Error Prediction in a Large Scale HPC System”, *48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks,* 2018, pp. 1–12.
16. J. Xiao et al., “Disk failure prediction in data centers via Online learning”, *Proceedings of the 47th International Conference on Parallel Processing,* 2018, pp. 1–10.
17. Y. Xu et al., “Improving Service Availability of Cloud Systems by Predicting Disk Error”, *2018 USENIX Annual Technical Conference,* 2018, pp. 481–494.
18. J. Alter et al., “SSD failures in the field: symptoms, causes, and prediction models”, *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2019, pp. 1–14.
19. Z. Shen et al., “Fast predictive repair in erasure-coded storage”, *49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, 2019, pp. 556–567.
20. J. Wang et al., “An Attention-augmented Deep Architecture for Hard Drive Status Monitoring in Large-scale Storage Systems”, *ACM Transactions on Storage* 15.3 (2019), pp. 1–26.
21. E. Xu et al., “Lessons and actions: What we learned from 10k ssd-related storage system failures”, *2019 USENIX Annual Technical Conference*, 2019, pp. 961–976.
22. J. Yu, “Hard disk Drive Failure Prediction Challenges in Machine Learning for Multi-variate Time Series”, *Proceedings of the 2019 3rd International Conference on Advances in Image Processing*, 2019, pp. 144–148.
23. S. Han et al., “Toward Adaptive Disk Failure Prediction via Stream Mining”. *2019 IEEE 40th International Conference on Distributed Computing Systems*, 2020.
24. S. Lu et al., “Making Disk Failure Predictions SMARTer!” *18th USENIX Conference on File and Storage Technologies*, 2020, pp. 151–167.
25. Backblaze. *Hard Drive Data and Stats*. https://www.backblaze.com/b2/hard-drive-test-data.html. [Online; accessed on 2020-04-28].

BIOGRAPHIES

Riccardo Pinciroli
Gran Sasso Science Institute
L’Aquila, 67100, Italy

riccardo.pinciroli@gssi.it

Riccardo Pinciroli received his M.S. (2014) and Ph.D. (2018) degrees from Politecnico di Milano. He is a Postdoc Fellow at the Gran Sasso Science Institute. His research interests include stochastic modeling, performance evaluation, reliability analysis, and uncertainty propagation.

Lishan Yang
Computer Science
William & Mary
Williamsburg, VA 23187-8795 USA

lyang11@email.wm.edu

Lishan Yang is a Ph.D. candidate in Computer Science at William & Mary. She received her B.S. degree in Computer Science from University of Science and Technology of China (USTC) in 2016. Her research interests are in GPU architecture, reliability analysis, and performance analysis.

Jacob Alter
Ipsos North America

jralter@email.wm.edu

Jacob Alter received an M.S. in Computer Science from

Authorized licensed use limited to: Gran Sasso Science Institute. Downloaded on April 18,2023 at 10:13:49 UTC from IEEE Xplore. Restrictions apply.

William & Mary in 2019. He is currently a Data Scientist at Ipsos North America. His research interests include storage systems, reliability analysis, and the application of data mining and machine learning to computing systems.

Evgenia Smirni
Computer Science
William & Mary
Williamsburg, VA 23187-8795 USA
esmirni@cs.wm.edu

Evgenia Smirni is the Sidney P. Chockley professor of computer science at William & Mary. Her research interests include queuing networks, stochastic modeling, resource allocation, storage systems, cloud computing, workload characterization, and modeling of distributed systems and applications. She is an ACM Distinguished scientist and an IEEE Fellow.

Authorized licensed use limited to: Gran Sasso Science Institute. Downloaded on April 18,2023 at 10:13:49 UTC from IEEE Xplore. Restrictions apply.