

Failure Prediction of Hard Disk Drives in Redundant Arrays Using Disk-Level Performance Metrics

Masanao Natsumeda¹

¹ *NEC Corporation, Kawasaki, Kanagawa, 211-8666, Japan
mnatsumeda@nec.com*

ABSTRACT

Although many hard drive failure prediction methods utilize Self-Monitoring Analysis and Reporting Technology (SMART) features, they are not collected in IT systems with demanding performance requirements to meet their specification. We present a novel data-driven method for the prediction utilizing disk-level performance metrics collected by Redundant Array of Independent Disk (RAID) controllers instead of SMART features. The proposed method computes relational anomaly scores leveraging logical relationships of Hard Disk Drives (HDDs) based on RAID configuration for better failure prediction. In addition, it further utilizes error codes from HDDs to filter out false positives. We evaluate the proposed method on a real-world dataset collected for this study from 881 disks used in disk arrays of RAID-6 and 1660 disks used in disk arrays of RAID-10 in a data center. The results show consistent performance improvement by the logical relationships and error-code-based filtering. In addition, seven out of nine failures are predicted one day before the failure at the latest. This result suggests that the proposed method provides plenty of time for HDD replacement before a failure occurs.

1. INTRODUCTION

A massive amount of data is generated in various systems. The total amount of new data generated in 2025 is forecasted to reach 175 ZB, which is five times larger than that in 2018 (Gantz, Reinsel, & Rydning, 2019). Storage devices are the key component to accommodate the large volume of data, and Hard Disk Drives (HDDs) will remain as the primary storage devices because of cost advantage.

In order to prevent data loss, a disk array consisting of multiple HDDs is virtualized with Redundant Array of Independent Disk (RAID) in enterprise IT systems. There are different levels of RAID to fulfill different needs of storage sys-

tems. For example, RAID-10 yields higher durability than RAID-6, but RAID-10 results in lower capacity than RAID-6. Although RAID typically increases durability, scenarios resulting in data loss remain. If the number of failed HDDs exceeds the limit for recovery, data loss occurs. Due to data security concerns, when a failure occurs in a disk array, other still functioning HDDs in the same disk array are replaced with new ones as preventive measures.

HDD failure prediction is a promising technology to reduce the running costs of IT systems and environmental load by the preventive measures since it allows us to replace HDDs only when necessary. Most HDD failure prediction methods (Aussel et al., 2017; Featherstun & Fulp, 2010; Ganguly et al., 2016; Hamerly, Elkan, et al., 2001; Li et al., 2014; Murray, Hughes, & Kreutz-Delgado, 2003; Yang et al., 2020; Zhang, Huang, Zhou, Xie, & Schelter, 2020) rely on Self-Monitoring Analysis and Reporting Technology (SMART) features. However, due to demanding performance requirements, SMART features are not always available in enterprise IT systems. On the other hand, two types of data representing HDD's status, i.e., disk-level performance metrics collected by RAID controllers and error codes from HDDs, do not require additional workload on disks for the collection. Therefore, HDD failure prediction utilizing them is applicable to a wide range of enterprise IT systems. An example of the error code sequence from an HDD is shown in Table 1. It shows error codes were issued before the HDD failed, indicating they capture disk degradation.

HDDs in the same disk array have logical relationships since they are virtualized as a single logical unit. The relationships appear as similar patterns in some of the disk-level performance metrics. For example, either RAID-1 or RAID-10 simultaneously write a pair of HDDs, and that leads similar pattern in their metrics related to writing data. Time series of average write response time, one of the disk-level performance metrics, from two HDDs in the same disk array are shown in Fig. 1. The HDD of the upper chart failed at the right-most of the chart, and the other remains normal. Both of the series similarly increase at the same time, but the in-

Masanao Natsumeda et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Table 1. An example of the error code sequence from an HDD.

Time	Message
2021-05-29 03:38:09	Unrecovered read error
2021-05-29 03:38:31	Unrecovered read error
2021-05-29 03:38:38	Unrecovered read error
2021-05-29 03:38:41	Data access failures occurred frequently and exceeded the threshold.
2021-05-29 03:38:41	The failed drive was removed from the RAID array.

crease of the failed HDD at the nearest time of the failure is significantly larger than that of the normal HDD. Due to their similar patterns during the healthy condition, their difference around the failure more emphasizes the abnormal increase than their individual metrics. This observation implies that the logical relationships by RAID configuration improves the accuracy of failure prediction.

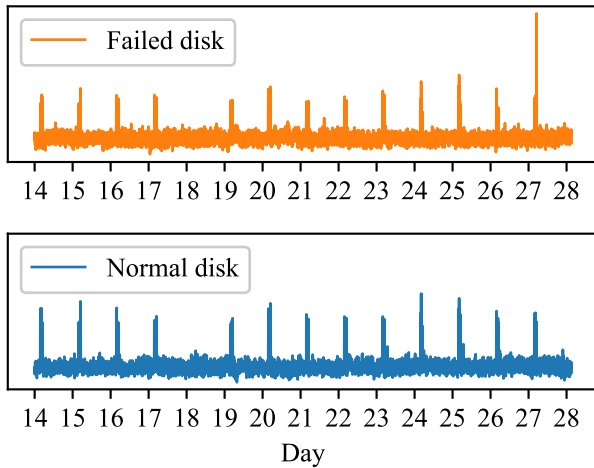


Figure 1. The motivated example. Both charts show the average write response time of HDDs in the same disk array. The scales of both charts are the same. The right-most time is of the HDD failure. They usually synchronize with each other, but the failed disk shows a higher value near the failure time.

In order to predict HDD failures without SMART features, this study proposes a two-step unsupervised anomaly detection method utilizing the logical relationships of disk-level performance metrics collected by RAID controllers and error codes from HDDs. The proposed method computes anomaly scores for each HDD. Then it adjusts the scores in the same disk array based on themselves, suppressing false positives because of sudden changes in disk-level performance metrics. In addition, the proposed method only detects an anomaly in

an HDD when error codes from the HDD are issued in advance. This error-code-based filtering further reduces false positives since the error codes provide conclusive evidence for disk degradation. When the proposed method detects an anomaly in an HDD, it is regarded as an early indication of a failure, and the replacement of the HDD is recommended.

2. RELATED WORK

HDD failure prediction has been studied for decades, and the studies mainly rely on SMART features. Hamerly et al. (2001) proposed two Bayesian approaches to the prediction, viewing the problem as anomaly detection. Murray et al. (2003) compared the performance of Support Vector Machines (SVMs), unsupervised clustering, and non-parametric statistical test. Li et al. (2014) employed Classification and Regression Trees for accurate, stable, and interpretable prediction for HDD failure. It was evaluated on a real-world dataset containing 25,792 drives. Zhang et al. (2020) addressed unsatisfactory results because of a small amount of data during training or the disks which have not appeared in training, employing a Long Short-Term Memory (LSTM) based siamese network. Featherstun and Fulp (2010) used all the data collected by syslog, a standard Unix logging facility. The data contains SMART features as well. Ganguly et al. (2016) combined multiple data sources, i.e., SMART features and Windows performance counters. Aussel et al. (2017) evaluated machine learning models on a large-scale and heterogeneous dataset from over 47,000 HDDs with 81 models from 5 manufacturers. Lu et al. (2020) performed large-scale disk failure analysis based on 380,000 HDDs distributed across 64 data center sites. They collected SMART features, disk-level performance metrics, server-level performance metrics, and disk spatial location data. The analysis was conducted with supervised methods for HDD failure prediction. The result shows that the supervised methods yield competitive performance without SMART features. Unlike the above studies, this study explores the feasibility of unsupervised HDD failure prediction without SMART features. In addition, this study utilizes error codes and logical relationships instead of physical relationships, i.e., location, for better prediction.

3. METHODOLOGY

The proposed method consists of RAID-configuration-based anomaly detection and error-code-based filtering. Its entire pipeline is shown in Fig. 2. RAID-configuration-based anomaly detection computes relational anomaly scores indicating the abnormality and detects an anomaly when it exceeds its pre-determined threshold. When an anomaly in an HDD is detected, the error-code-based filtering checks if any error code has been issued to the HDD before. If it has been issued, the detection is regarded as negative, and the proposed method continues monitoring the HDD.

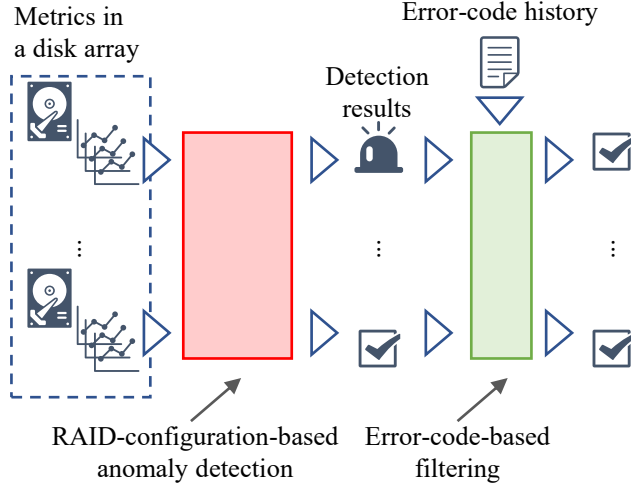


Figure 2. The entire pipeline of the proposed method. Detection results by RAID-configuration-based anomaly detection is updated by error-code-based filtering.

The pipeline computing relational anomaly scores is shown in Fig. 3. It consists of two functions, i.e., the individual scoring function and the relational scoring function. The individual scoring function computes a score based on disk-level performance metrics for each HDD. The score called individual score represents the abnormality of the HDD at the time stamp associated with the metrics when logical relationships between HDDs in the same disk array are not incorporated. The relational scoring function adjusts individual scores associated with the same time stamp in the same disk array. The adjusted score is the relational anomaly score.

The individual scoring function is realized with the distance-based method called k NN (Angiulli & Pizzuti, 2002). k NN assigns an anomaly score by taking the average distance over K -nearest-neighbors in the training set. Let $x' \in \mathbb{R}^d$ be d dimensional metrics from an HDD at testing and $x^k \in \mathbb{R}^d$ be k th nearest neighbor of x' in the training set. The individual score a'_1 is defined as:

$$a'_1 = \frac{1}{K} \sum_{i=1}^k \|x' - x^i\|_2^2. \quad (1)$$

Since the main concern of this study is whether the logical relationships of RAID help improve HDD failure prediction, k NN is used due to its simplicity; however any unsupervised anomaly detection method is applicable to the individual scoring function.

The relational scoring function computes distances of the individual scores from their mean at the same time stamp as relational anomaly scores. Given a disk array, let N be the number of HDDs in the disk array and a'_{1n} be the individual score of the n th HDD. The relational anomaly score a'_{Rn} is

defined as:

$$a'_{Rn} = a'_{1n} - \frac{1}{N} \sum_{i=1}^n a'_{1i}. \quad (2)$$

4. EXPERIMENTS

We evaluate the proposed method on a private dataset and compare it to its three variants to examine the benefits of adjusting individual scores based on RAID configuration and error-code-based filtering. The first variant does not employ error-code-based filtering. The second variant only relies on individual scores by k NN. The third variant only relies on error codes. The first two variants predict HDD failure based on anomaly scores. On the other hand, the third variant predicts HDD failure when an error code is issued. The third variant does not consider the category and occurrence frequency of the codes for prediction since the proposed method does not consider them for filtering.

4.1. Dataset

The dataset consists of time series of disk-level performance metrics and error code sequences of HDDs used in a data center. All the HDDs are Small Computer System Interface (SCSI) devices. The disk-level performance metrics used in this experiment are shown in Table 2. The total number of metrics is 10. Their readings were collected every three minutes. Training data consists of time series of the metrics from 19 normal HDDs. The HDDs were sampled from different disk arrays of RAID-6. The data collection period for each HDD ranges from 13 to 15 days. Testing data consists of time series of the metrics and error code sequence from 43 disk arrays of RAID-6 and 106 disk arrays of RAID-10. Since HDD failures rarely occur, the testing data collected for a longer period than that of the training data. The data collection period is 4.5 months. The total number of HDDs used in the disk arrays of RAID-6 and RAID-10 are 862 and 1660, respectively. The number of failed HDDs in the disk arrays of RAID-6 and RAID-10 are two and seven, respectively.

Table 2. The disk-level performance metrics used in this experiment.

Index	Name
1	Write Count
2	Write Average Response Time
3	Write Max Response Time
4	Write Transfer Rate
5	Read Count
6	Read Average Response Time
7	Read Max Response Time
8	Read Transfer Rate
9	Busy Ratio
10	Busy Time

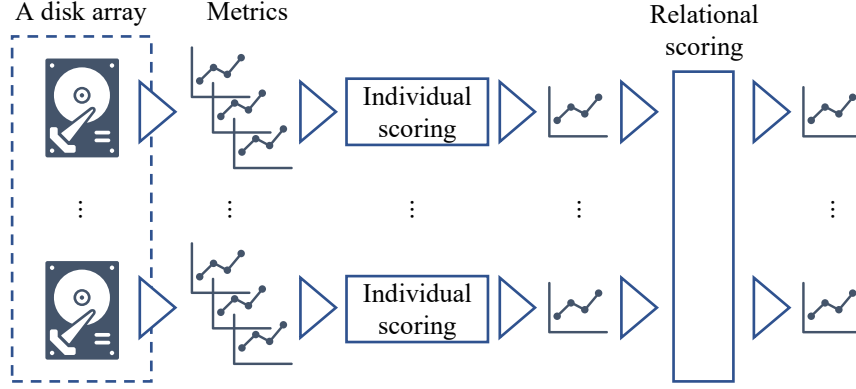


Figure 3. The pipeline computing relational anomaly scores.

4.2. Experimental setting

The HDD prediction is evaluated segment basis rather than time-stamp basis. The time series of the metrics are divided into time series segments at the time of failure as shown in Fig. 4. Each segment is regarded as a sample for evaluation. The number of nearest neighbors K is set to 10. Area Un-




Figure 4. Sample generation for the segment-basis evaluation. The horizontal axis shows the entire period of testing data collection. Each time series segment divided by the time of failure forms a sample. This example generates four samples.

der the Curve (AUC), precision, recall, and F1-score are used as evaluation metrics. Let TP be the number of true positive, FP be the number of false positive, TN be the number of true negative, FN be the number of false negative, True positive rate TPR is defined by:

$$TPR = \frac{TP}{TP + FN}. \quad (3)$$

It is the same as recall r . False positive rate FPR is defined by:

$$FPR = \frac{FP}{FP + TN}. \quad (4)$$

Precision p is defined by:

$$p = \frac{TP}{TP + FP}. \quad (5)$$

F1-score f_1 is defined as:

$$f_1 = \frac{2 \times r \times p}{r + p}. \quad (6)$$

AUC is the area underneath the entire Receiver Operating Characteristic (ROC) curve. Suppose TPR is a function of FPR . AUC is computed as:

$$\int_0^1 TPR(FPR) dFPR. \quad (7)$$

The same value is used as the threshold for each of the HDDs. Due to the small number of failures, it is difficult to determine an appropriate threshold. Thus we report the theoretical best values of the evaluation metrics, adjusting the threshold for each evaluation metric on the testing data. Remaining time from the detection time until the failure is also evaluated for each failed HDD. For this evaluation, the threshold is set to give the maximum precision subject to 100% of recall.

4.3. Results

Performance comparison of different methods of computing anomaly scores is shown in Table 3. The proposed method is denoted by ‘relational + log’. Its variant without error-code-based filtering is denoted by ‘relational’. The method only relying on individual scores by k NN is denoted by ‘individual’. The proposed method shows the best results for all the evaluation metrics. The performance improvement by the unique components in the proposed method is consistent over different evaluation metrics. The proposed method performs better than its variant without error-code-based filtering. The variant without error-code-based filtering performs better than the method relying only on individual scores.

To better illustrate the consistent performance improvement, ROC curves of different methods on RAID-10 and RAID-6 are shown in Fig. 5 and Fig. 6, respectively. The consistent performance improvement is attributed to a consistently

Table 3. Performance comparison of different methods of computing anomaly scores. P, R, and F1 represent precision, recall, F1-score, respectively. The bold letters indicate the best results.

Method	RAID-6				RAID-10			
	AUC	P	R	F1	AUC	P	R	F1
individual	0.989	0.500	1.000	0.500	0.895	0.075	1.000	0.128
relational	0.990	0.500	1.000	0.500	0.950	0.077	1.000	0.132
relational + log	0.998	1.000	1.000	0.667	0.990	0.214	1.000	0.343

higher true positive rate at the same false positive rate. For any false positive rate, the proposed method consistently outperforms its variant without error-code-based filtering, and the variant consistently outperforms the method only relying on individual scores.

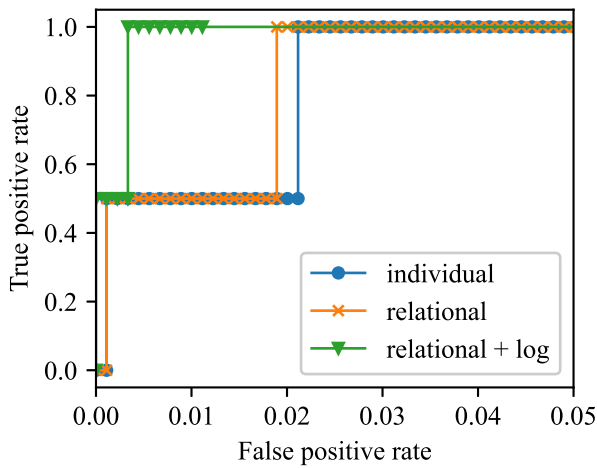


Figure 5. ROC curves on RAID-6.

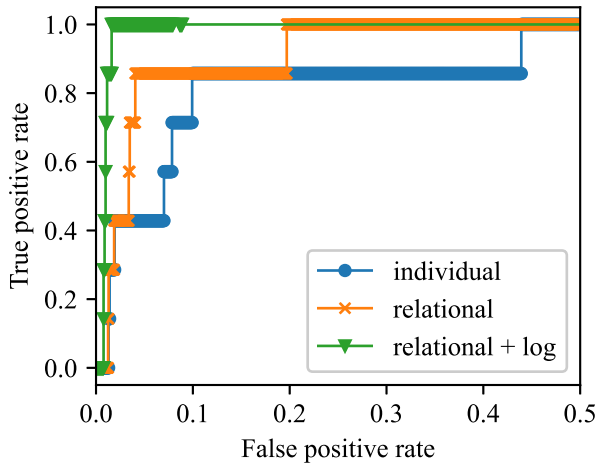


Figure 6. ROC curves on RAID-10.

The performance comparison of the proposed method and its variant only relying on error codes is shown in Table 4. This table also indicates the effect of disk-level performance metrics on failure prediction. Since error codes appear before

failures for all the cases, precision at 100% recall is compared. The proposed method shows the best results for both RAID levels, and the precision is more than twice that of the variant, indicating the positive effect of the disk-level performance metrics.

Table 4. The effect of disk-level performance metrics. Precision at 100% recall is compared. The bold letters indicate the best results.

Method	Precision	
	RAID-6	RAID-10
log	0.167	0.080
metrics + log	0.400	0.184

Distribution of remaining time from the detection time until the failure is shown in Table 5. Seven out of nine failures are predicted one day before the failure at the latest. This result suggests that the proposed method provides plenty of time for HDD replacement before a failure occurs.

Table 5. Distribution of remaining time from the detection time until the failure.

Remaining time	RAID-6	RAID-10	Total
~ 1 month	2	2	4
1 month ~ 2 weeks	0	1	1
2 weeks ~ 1 week	0	0	0
1 week ~ 3 days	0	1	1
3 days ~ 1 day	0	1	1
1 days ~ 12 hours	0	2	2

5. CONCLUSION

This study has proposed a two-step unsupervised anomaly detection method utilizing the logical relationships of disk-level performance metrics collected by RAID controllers and error codes from disks instead of SMART features. Evaluation of the proposed method on the private dataset has verified the benefits of incorporating RAID configuration and error codes from HDD into HDD failure prediction. The proposed method leaves plenty of room for improvement. For example, a more complex method is applied to compute individual scores instead of k NN, and the error-codes-based filtering incorporates the category and occurrence frequency of the codes into the prediction. Although further evaluation of the

proposed method with a more significant amount of data is necessary due to the small volume of data for evaluation, we hope that our observations from experiments facilitate a new research direction of HDD failure prediction, not relying on SMART features.

REFERENCES

- Angiulli, F., & Pizzuti, C. (2002). Fast outlier detection in high dimensional spaces. In *Principles of data mining and knowledge discovery: 6th european conference, pkdd 2002 helsinki, finland, august 19–23, 2002 proceedings 6* (pp. 15–27).
- Aussel, N., Jaulin, S., Gandon, G., Petetin, Y., Fazli, E., & Chabridon, S. (2017). Predictive models of hard drive failures based on operational data. In *2017 16th IEEE international conference on machine learning and applications (icmla)* (pp. 619–625).
- Featherstun, R. W., & Fulp, E. W. (2010). Using syslog message sequences for predicting disk failures. In *Lisa*.
- Ganguly, S., Consul, A., Khan, A., Bussone, B., Richards, J., & Miguel, A. (2016). A practical approach to hard disk failure prediction in cloud platforms: Big data model for failure management in datacenters. In *2016 IEEE second international conference on big data computing service and applications (bigdataservice)* (pp. 105–116).
- Gantz, J. F., Reinsel, D., & Rydning, J. (2019). *The us data-sphere: Consumers flocking to cloud*.
- Hamerly, G., Elkan, C., et al. (2001). Bayesian approaches to failure prediction for disk drives. In *Icml* (Vol. 1, pp. 202–209).
- Li, J., Ji, X., Jia, Y., Zhu, B., Wang, G., Li, Z., & Liu, X. (2014). Hard drive failure prediction using classification and regression trees. In *2014 44th annual IEEE/IFIP international conference on dependable systems and networks* (pp. 383–394).
- Lu, S., Luo, B., Patel, T., Yao, Y., Tiwari, D., & Shi, W. (2020). Making disk failure predictions smarter! In *Fast* (pp. 151–167).
- Murray, J. F., Hughes, G. F., & Kreutz-Delgado, K. (2003). Hard drive failure prediction using non-parametric statistical methods. In *Proceedings of icann/iconip*.
- Yang, Q., Jia, X., Li, X., Feng, J., Li, W., & Lee, J. (2020). Evaluating feature selection and anomaly detection methods of hard drive failure prediction. *IEEE Transactions on Reliability*, 70(2), 749–760.
- Zhang, J., Huang, P., Zhou, K., Xie, M., & Schelter, S. (2020). Hddse: Enabling high-dimensional disk state embedding for generic failure detection system of heterogeneous disks in large data centers. In *Proceedings of the 2020 usenix conference on usenix annual technical conference* (pp. 111–126).