

Anomaly detection using SMART indicators for hard disk drive failure prediction

Sladana Đurašević, Borislav Đorđević

Abstract— Hard disk drives are more prone to failures since they are the only electromechanical component of computer systems susceptible to mechanical wear. Failure of hard disk drive leads to permanent data loss, which is typically more costly than drive itself. Prediction of hard drive failure enables user notification to copy stored data to another storage device, preventing data loss. SMART technology monitors vital hard disk drive's parameters and warn user when some of them exceed manufacturer defined threshold. Advanced failure prediction algorithms rely on machine learning to find mutual dependence of SMART indicators in order to provide more precise prediction of hard drive failure. In this paper, we presented algorithm based on anomaly detection method which enabled prediction of hard drive failure, 39 days prior actual failure of the hard drive.

Index Terms— Hard disk drive; Failure prediction; SMART; Machine learning; Anomaly detection

I. INTRODUCTION

Hard disk drives (HDDs) are still the only electromechanical component of the computer system which is due to its mechanical design more prone to failure than other components. HDD failure could become costly to user due permanent data loss, since they are used as primary storage of user data. HDD failures can be classified into predictable and unpredictable failures [1]. Predictable failures are caused by processes which slowly degrade drive performance due mechanical wear and gradual degradation of storage surfaces. Indicators of these process such problems with reading and writing of data, increase in number of damaged disk sectors, increased vibration level, can be monitored to determine when such failures are becoming more likely. Unpredictable failures represent sudden drive failures, which occur due defective electronic components or sudden mechanical failures caused by improper handling. Self-Monitoring, Analysis and Reporting Technology, known as SMART, is used to monitor various indicators of HDD operation. These indicators store information about drive temperature, operating hours, the number of on/off cycles, the number of damaged sectors and are used to indicate a possible imminent drive failure. HDD manufactures commonly define thresholds for each monitored SMART indicator, thus the user is notified about possible drive failure when certain SMART

indicator exceeds the predefined threshold.

A field study [2] conducted on 100,000 consumer-grade HDDs found correlations between certain SMART indicators and actual failure rates. It also shown that 36 % of drives failed without recording changes in any of SMART indicators, which limits usefulness in anticipating failures. Authors of papers [3] derived distribution-free statistical hypothesis tests which improve failure prediction. More advanced failure prediction algorithms exploit mutual dependence of multiple SMART indicators to make a more accurate prediction of drive failure [4, 5]. Such models rely on data sets of SMART indicators collected from the large population of hard drives, operating under similar conditions. Such data sets are collected in data centers and are typically inaccessible to researchers. In this paper we applied generic anomaly detection method to create anomaly detection algorithm for prediction of HDD failure. Derived algorithm is trained using data set of SMART indicators from large population of HDDs and results on independent data set shown high precision of failure prediction.

II. SMART INDICATORS

Backblaze, a remote backup service company, started sharing SMART statistics of the HDDs operating in their data center since 2013 [6]. Among different HDD models used in their data center, Seagate ST3000DM001 stands out with the substantial percentage of failures. This 3 TB hard disk drive produced by Seagate Technology from 2011, uses three 1 TB platters rotating at the spindle speed of 7200 rpm. This drive is intended to be used in desktop systems, direct-attached external storage devices (DAS) and network-attached storage devices (NAS). Manufacturer rates this HDD at 300 000 load/unload cycles with annular failure rate (AFR) of less than 1% per 2400 hours of operation per year. AFR is a percentage estimate of the products that will likely fail due to a defect over a 1-year period operated at nominal use level. Nominal use level for desktop HDDs is 2400 hours per year which corresponds to drive operation 8 hours daily five days per week. In case enterprise HDD used in servers, nominal use level of 8760 hours corresponds to continuous 24/7 operation all around the year. Starting from the population of 4829 ST3000DM001 HDDs operated by Backblaze from 2012, 1880 drives failed so far as shown in Fig 1. Furthermore, about 75% of reminder drives failed on of the bench tests once they were removed from service. Such high failure rates prompted a class action against Seagate to be filed in 2016, and primarily cited reliability data provided by Backblaze.

Sladana Đurašević is with the Faculty of Technical Sciences in Čačak, University of Kragujevac, Svetog Save 65, 32000 Čačak, Serbia (e-mail: sladjana.djurasevic@gmail.com).

Borislav Đorđević is with the Institute Mihailo Pupin, Volgina 15, 11060 Beograd, Serbia (e-mail: bora@impcomputers.com).

Some experts argue that Backblaze used the drives in a manner that "far exceeded the warranty conditions", like 24/7 operation of drive intended for 8 hour per day desktop use.

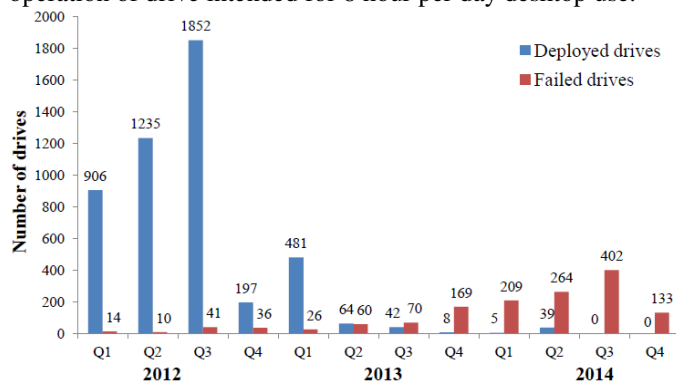


Fig. 1. Number of deployed and failed ST3000DM001 drives (2012 ÷ 2014)

SMART indicators in Backblaze data set are available in raw and normalized form. The raw form of SMART indicators contains decimal values representing as a number of damaged sectors, operating hours or temperature, but sometimes raw parameters don't provide any meaningful information as a decimal number. The normalized form of SMART indicators is obtained by scaling of raw parameter into the range between 0 and 100, according to vendor specific scale. Backblaze data set is collected daily for each operational hard drive. Each entry contains the date of entry, basic drive information, failure indication and the set of SMART indicators. Basic drive information contains model, capacity and the manufacturer assigned serial number of the drive. Failure indication contains a "0" for the operational drive and is set to "1" if this is the last day the drive was operational before failing. Set of SMART indicators contains raw and normalized SMART data for up to 45 different SMART parameters depending on drive model.

III. ANOMALY DETECTION METHOD

Anomalies represent set of events which occur relatively infrequently are considerably dissimilar from the remainder of the data. When anomalies occur, their consequences can be quite dramatic and quite often in a negative sense [7]. Anomalies can be detected using statistical-based, distance-based or model-based methods. Statistical methods rely on the statistical representation of "normal" behavior of overall population for input data vector x , the most commonly with normal distribution or some more appropriate distribution. Normal distribution $p(x)$ of normal behavior is then used to determine confidence limit ϵ for detecting anomalies using the training set for input vector x . Anomalies are observations whose characteristics differ significantly from the normal behavior. This training set contains a certain number of anomalies which are used to find the most appropriate confidence limit, $p(x) < \epsilon$, in order to correctly detect as many anomalies as possible [8]. In order to determine the most appropriate confidence limit ϵ , certain evaluation metrics, derived from confusion matrix should be used. The confusion matrix is used to represent the numbers of successful anomaly

detections, called True positives, the number of incorrect anomaly detections, called false positives, and the number of undetected anomalies, called false negatives.

Based on values from confusion matrix, we can use following metrics for determining the most appropriate confidence limit: Precision, Recall and F-score. Precision represents the ratio between the number of successfully detected anomalies and the total number of detected anomalies and is used to measure the accuracy of anomaly detection model. Accuracy metrics can be misleading because the model will tend to detect the small number of anomalies in order not to make inaccurate detections. The recall represents the ratio between the number of successfully detected anomalies and the total number of anomalies present in the dataset and is used to measure the percentage of detected anomalies. Recall metrics can be misleading because the model will tend to detect the huge number of anomalies in order detect all anomalies. In order to create the balance between precision and recall metrics F-score metrics represent the hybrid solution between precision and recall metrics.

IV. EXPERIMENTAL RESULTS

Anomalies occur when certain SMART parameters deviate from normal values, which usually lead to disk drive failure. We applied generic anomaly detection method based on statistical distribution [8] to create algorithm for drive failure prediction based on the most critical SMART. In order to achieve high accuracy of failure prediction, we choose precision as a metric for determining confidence limit.

Anomaly detection algorithm is applied on set of SMART indicators for Seagate ST3000DM001 disk drive. Data set contained entries for 4255 ST3000DM001 drives which operated for 1081649 hours from February 2014 to November 2015, from which 1357 have failed during regular operation. The dataset contains 24 SMART indicators represented in both raw and normalized form shown in Table I, among which the ones most likely to indicate failure were used for failure prediction.

SMART 1 (Read Error Rate) indicator represents the rate of hardware read errors that occurred during reading data from a disk surface. SMART 5 (Reallocated Sectors Count) indicator represents a number of bad sectors that have been remapped to spare area. SMART 7 (Seek Error Rate) indicator represents seek errors of magnetic heads caused by the partial failure in the mechanical positioning system. SMART 183 (Runtime Bad Block) represents a total number of uncorrectable errors encountered during normal operation. SMART 187 (Reported Uncorrectable Errors) represents a number of errors that could not be recovered by hardware error code correction. SMART 189 (High Fly Writes) represents a number of write operations performed when a recording head is flying outside its normal operating range. SMART 193 (Load Cycle Count) represents a number of cycles when magnetic heads are put into head landing zone position, as a result of power saving. SMART 197 (Current Pending Sector Count) represent the number of unstable sectors which are waiting to be remapped.

TABLE I
LIST OF MONITORED SMART INDICATORS

SMART indicator	Apr 2013 Jan 2014		Feb 2014 Nov 2015	
	Raw	Norm	Raw	Norm
SMART 1 - Read Error Rate	+		+	+
SMART 3 - Spin Up Time			+	+
SMART 4 - Start Stop Count			+	+
SMART 5 - Reallocated Sector Count	+		+	+
SMART 7 - Seek Error Rate			+	+
SMART 9 - Power On Hours	+		+	+
SMART 10 - Spin Retry Count			+	+
SMART 12 - Power Cycle Count			+	+
SMART 183 - Runtime Bad Block			+	+
SMART 184 - End to End Error			+	+
SMART 187 - Reported Uncorrected			+	+
SMART 188 - Command Timeout			+	+
SMART 189 - High Fly Writes			+	+
SMART 190 - Airflow Temperature			+	+
SMART 191 - G-Sense Error Rate			+	+
SMART 192 - Power-off Retract Cnt			+	+
SMART 193 - Load Cycle Count			+	+
SMART 194 - Drive Temperature	+		+	+
SMART 197 - Current Pending Sector	+		+	+
SMART 198 - Offline Uncorrectable			+	+
SMART 199 - UDMA CRC Error Cnt			+	+
SMART 240 - Multi Zone Error Rate			+	+
SMART 241 - Total LBAs Written			+	+
SMART 241 - Total LBAs Read			+	+

Dataset is separated into two subsets, first containing 2989 healthy drives and second containing 1357 failed drives. The training set is composed of 60 % population of healthy drives in order to determine statistical distributions of important smart parameters, which are shown in Table II. Training set contained data from 1739 healthy drives with 516483 entries.

TABLE II
NORMAL DISTRIBUTIONS OF IMPORTANT SMART INDICATORS

SMART Indicator	Mean μ	Standard deviation σ
SMART 1	115.29	3.89
SMART 5	99.86	1.83
SMART 7	83.72	6.66
SMART 183	96.80	15.87
SMART 187	98.26	8.46
SMART 189	95.95	14.28
SMART 193	51.13	36.39
SMART 197	99.999	0.57

The confidence limit ϵ for anomaly detection was determined using cross-validation set, composed from data of 20 % of healthy drives and data of 50 % of failed drives. It contained data from 579 healthy drives and 679 failed drives with a total number of entries 291407. Confidence limit for selected SMART indicators was varied in range ($10^{-60} \div 10^{-10}$) were most suitable metrics chosen for anomaly detection was precision. The results show that very high precision of anomaly detection was achieved, as high as 0.936.

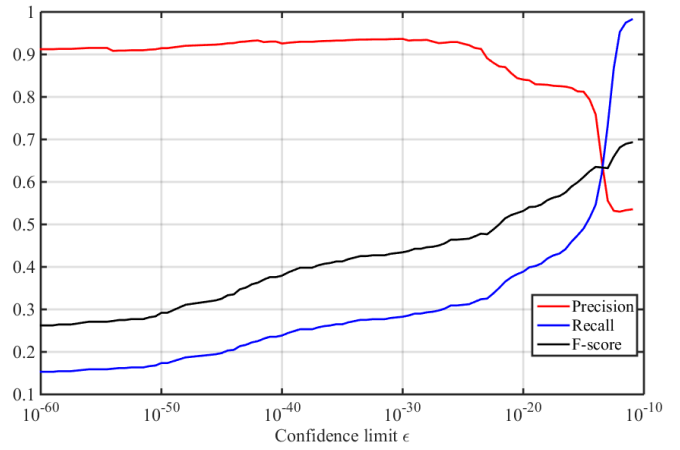


Fig. 2. Precision, Recall and F-score for various value of confidence limit

In order to measure the real accuracy of proposed anomaly detection model, the independent test set was composed of data from remaining 20 % of healthy drives and remaining of 50 % of failed drives. It contained data from 580 healthy drives and 678 failed drives with a total number of entries 273759. Selected Confidence for highest precision in cross-validation set was used to detect anomalies in the test set and results are presented in Table III. Results show that high precision of 0.911 was kept with the test set, with the recall of 0.304. The number of detected anomalies was 206 of 678 with only 20 false anomalies detected. Furthermore, some of the SMART indicators were discarded from the model in order to further improve the accuracy of anomaly detection. When indicators SMART 1 and SMART 7 were omitted, 210 of 678 failures were detected with 19 false alarms. Furthermore, using just most important indicators SMART 5, SMART 187 and SMART 197, we were able to increase the precision with 230 of 678 failures detected with 20 false alarms.

TABLE III
NORMAL DISTRIBUTIONS OF IMPORTANT SMART INDICATORS

Used SMART indicators	SMART 1 SMART 5 SMART 7 SMART 183 SMART 187 SMART 189 SMART 193 SMART 197	SMART 5 SMART 183 SMART 187 SMART 193 SMART 197	SMART 5 SMART 187 SMART 197
Cross validation			
Precision	0.936	0.938	0.940
Recall	0.283	0.287	0.325
F-score	0.434	0.439	0.484
ϵ	$1 \cdot 10^{-30}$	$1 \cdot 10^{-24}$	$5 \cdot 10^{-14}$
Test set			
Precision	0.911	0.917	0.920
Recall	0.304	0.310	0.339
F-score	0.456	0.463	0.496
True positive	206	210	230
False negative	472	468	448
False positive	20	19	20

V. CONCLUSION

Developed anomaly detection model was further analyzed in order to determine the time lag between failure prediction and actual failure of the drive. One example of the lifetime operation of the drive with serial number W1F08JSX was shown in Figure 3. Indicators SMART 5 and SMART 187 started to decline and anomaly detection model predicted failure after 97 operating days. The actual failure occurred after 154 operating days. Furthermore, we analyzed time taken between failure prediction and actual failure for all failed drive which is presented by the histogram in Figure 4. Results show that anomaly detection algorithm provided on average 38.9 days warning prior actual failure of HDD.



Fig. 3. Example of failure prediction for drive with serial number W1F08JSX




Fig. 4. Histogram of time between failure prediction and actual drive failure

The possibility of predicting failures of hard disk drives enables the user to take preemptive actions in order to backup important data. In this paper, we presented anomaly detection algorithm which is capable of predicting failure of the hard disk drive using SMART indicators. Presented model achieved high precision of failure detection over 90 %. Proposed anomaly detection model provided on average 38.9 days warning prior the actual failure of hard disk drive.

REFERENCES

- [1] B. Schroeder, G. 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 (FAST 2007), San Jose, USA, February 13-16, 2007, pp. 1-16
- [2] E. Pinheiro, W.D. Weber, L. A. Barroso, "Failure Trends in a Large Disk Drive Population", 5th USENIX Conference on File and Storage Technologies (FAST 2007), San Jose, USA, February 13-16, 2007, pp. 17-29
- [3] G. Hughes, J. Murray, K. K. Delgado, C. Elkan, "Improved Disk-Drive Failure Warnings", IEEE Transactions On Reliability, Vol. 51, No. 3, September 2002
- [4] J. Murray G. Hughes K.K. Delgado, "Machine Learning Methods for Predicting Failures in Hard Drives: A Multiple-Instance Application" Journal of Machine Learning Research 6 (2005), pp. 783-816
- [5] M. Botezatu, I. Giurgiu, J. Bogojeska, D. Wiesmann, Predicting Disk Replacement towards Reliable Data Centers, Conference of Knowledge discovery and data mining KDD2016, August 13 - 17, 2016, San Francisco, CA, USA
- [6] BackBlaze SMART data set, <https://www.backblaze.com/b2/hard-drive-test-data.html>
- [7] Y. Tan, X. Gu, "On Predictability of System Anomalies in Real World" IEEE International Symposium on Modeling, Analysis & Simulation of Computer and Telecommunication Systems (MASCOTS), 2010, 17-19 Aug. 2010, Miami Beach, FL, USA,
- [8] T. Dunning, Ted; E. Friedman, "Practical Machine Learning: A New Look at Anomaly Detection", O'Reilly publishing, 2014