A survivability analysis of enterprise hard drives incorporating the impact of workload

Aman Mallik[1]*, B Ranjith Reddy[2] and Gadadhar Sahoo[3]

[1]Department of Computer Science, BITs Pilani Hyderabad Campus, Hyderabad, India, [2]Msrit (MS Ramaiah Institute of Technology), Bangalore, India, [3]Department of Computer Science and Engineering, IIT (ISM) Dhanbad, Dhanbad, India

Introduction: Hard disk drive (HDD) failure is a significant cause of downtime in enterprise storage systems. Research suggests that data access rates strongly influence the survival probability of HDDs.

Methods: This paper proposes a model to estimate the probability of HDD failure, using factors such as the total data (TD) read or written and the average access rate (AAR) for a specific drive model. The study utilizes a dataset of HDD failures to analyze the effects of these variables.

Results: The model was validated using case studies, demonstrating a strong correlation between access rate management and reduced HDD failure risk. The results indicate that managing data access rates through improved throttle commands can significantly enhance drive reliability.

Discussion: Our approach suggests that optimizing throttle commands at the storage controller level can help mitigate the risk of HDD failure by controlling data access rates, thereby improving system longevity and reducing downtime in enterprise storage systems.

KEYWORDS

cox proportional hazards model, storage system risk reduction, HDD failure prediction, survivability probability, enterprise storage systems

1 Introduction

The fundamental role of any storage system is to efficiently cater to the diverse needs of various applications while accommodating their specific workloads. In modern storage systems, data reading and writing operations are primarily handled by hard disk drives (HDDs). Any interruption of these basic tasks can have significant repercussions, affecting all facets of storage management. Such disruptions can lead to performance degradation, an increased need for human intervention, a higher risk of service outages, and ultimately, potential data unavailability or loss.

Despite the critical importance of seamless data operations, many storage systems face challenges in maintaining uninterrupted performance. Existing solutions often fall short in addressing the complexity and variability of workloads, reducing the efficiency of data management. Existing research has focused on optimizing the individual components of storage systems, but comprehensive strategies that encompass the entire storage infrastructure are still lacking.

This research aims to bridge these gaps by exploring innovative approaches to enhancing the reliability and efficiency of storage systems. By investigating the underlying causes of data operation interruptions and their impact on overall system performance, this study seeks to develop robust solutions that minimize disruptions, reduce human intervention, and ensure continuous data availability. Through this work, we aim to contribute to the advancement of storage technologies, paving the way for more resilient and efficient data management systems.

1.1 Motivation

The motivation for this paper stems from the critical issue of HDD failure causing significant downtime in enterprise storage systems. Despite advancements in storage technology, HDD failures remain a persistent challenge, leading to substantial disruptions and operational inefficiencies. The existing literature provides valuable insights into individual factors that contribute to HDD failures, but there is a lack of comprehensive models that combine these factors to predict and mitigate such failures effectively.

In this work, empirical analyses are conducted to explore the relationship between data access rates and HDD survivability. By proposing a predictive model based on the total data (TD) read or written and the average access rate (AAR), this work aims to preemptively address HDD failures. The predictive model leverages real-world data from storage controllers to identify HDDs with high failure probabilities.

To validate the findings, the proposed model is tested using data from actual storage systems. A novel strategy is introduced: reallocating HDDs with high failure probabilities to different redundancy groups. This approach aims to optimize resource allocation, enhance system resilience, and mitigate the risks associated with HDD failures.

By addressing these challenges, this research provides a practical solution for storage system administrators and engineers and a proactive method to improve the reliability and efficiency of enterprise storage systems.

1.2 Previous research

HDDs consist of many complex subcomponents that must work in coordination with each other. Depending on the characteristics of the subcomponents, failures can occur at different stages in the lifetime of a product. The mean time to failure (MTTF) and annualized failure rate (AFR) are two of the current metrics of choice for quantifying the survivability of HDDs (George, 2013).

Prior research has shown that HDD failure prediction modeling can provide reasonable failure predictions for different kinds of hard disks with various interfaces, including integrated drive electronics (IDE), fiber channels (FC), small computer system interfaces (SCSI), and serial advanced technology attachments (SATA). Statistical modeling techniques such as logistic regression have been applied using the most relevant self-monitoring, analysis, and reporting technology (SMART) parameters to predict HDD failures with reasonable false alarm rates and accuracy (Shen et al., 2018; Zhang et al., 2023; Rincón et al., 2017; Liu and Xing, 2020; Smith and Smith, 2004; Smith and Smith, 2001; Smith and Smith, 2005; Mohanta and Ananthamurthy, 2006).

There have been attempts to increase the survival probability of solid-state driFDrieves (SSDs) in storage systems by controlling the write amplification factor (WAF) depending on the workload (Mohanta et al., 2015). Different machine-learning solutions have also been proposed as alternatives, including support vector machines, nonparametric rank-sum tests, and unsupervised clustering algorithms. These solutions provide improvements over existing threshold-based algorithms for predicting HDD failures (Murray et al., 2003; Hamerly and Elkan, 2001; Royston and Sauerbrei, 2008).

Improvements and new trends in HDD packaging methods have led to the high-density packing of physical materials in the drive cage. This has led to less spacing between the head and the media during read and write operations, which is one of the probable causes of HDD failure with media errors. The MTTF parameter alone is inadequate to describe HDD survivability, and according to one study (George, 2013), the total amount of data transferred is a more appropriate parameter of choice.

Many large enterprise storage companies require that drive manufacturers to provide detailed information to be queried from HDDs for further failure analysis. Along with the standard parameters, such as SMART attributes, recoverable errors, and unrecoverable errors parameter values are also made available to query from HDDs. This research employs these data for predictive modelling and analyses.

1.3 Unique contributions

Different kinds of hard disks form components of the storage systems deployed in data centers and cloud environments, including private, public, and hybrid clouds, to provide platform-as-a-service (PaaS) technology. In storage systems, these hard disks are building blocks for offering proper performance and survivability to the hosted application. Disk failure is inevitable, and catastrophic errors threaten both mission-critical data and the performance and survivability of applications.

The novel application of the model proposed in this work is that multiple HDDs in the same storage pool with a high failure probability threshold are reallocated to different storage pools after ensuring proper migration of their data to enhance overall system survivability.

The following are unique contributions of this work that advance the understanding of hard drive failures and their prediction:

- **Predictive Modeling:** By leveraging the statistical programming language R, the proposed predictive model has the capability to analyze workload patterns and their correlation with hard drive failures. By training models on extensive datasets that incorporate workload parameters, these algorithms can forecast potential failures more accurately.
- **Dynamic Adjustments:** The same model can be pragmatically implemented for developing real-time monitoring tools that continuously assess workload parameters and dynamically adjust hard drive operations. This proactive approach allows systems to mitigate potential stressors or redistribute the workload to ensure optimal drive health.
- **End-to-End Assessment:** This model has the potential to perform a comprehensive lifecycle analysis that considers the cumulative impact of workload variations over the entire lifespan of a hard drive. The appropriately chosen statistical model, namely, the Cox Proportional Hazards Model (CPHM) utilizes censored data to provide with valid estimates of the survival probabilities for potential HHD failures throughout the entire life cycle. This holistic approach helps in identifying critical periods or thresholds beyond which workload intensities significantly affect survivability.
- **Smart Resource Allocation**: This model unveils avenues for ushering in adaptive throttling mechanisms that optimize resource allocation based on workload analysis. By dynamically

adjusting read/write operations, applying caching strategies, and making use of data placement, systems can reduce wear and tear, thereby enhancing hard drive longevity.

- **Unified Metrics**: This model facilitates the use of proposing standardized metrics that incorporate workload considerations based on the AAR and TD. By establishing a common framework for evaluating hard drive survivability across diverse workloads, researchers and industry practitioners can compare results more effectively and drive advancements collaboratively.

1.4 Mathematical formulations and equations

Storage Array Downtime (SD) is defined as the percentage of time that the storage arrays in the installed base are down and not available to service requests. Measuring the storage array downtime is useful for clearly visualizing different contributions to the total array downtime (such as firmware-related downtime and downtime with other causes). However, the proposed research has adopted another approach is adopted in this work to investigate the impact of workload on HDD failures. As part of this analysis, primary data was collected from real production systems. The total amount of data read from or written to HDDs is a major parameter of this analysis. It is called the total data (TD) and is defined as follows:

$$TD = TDR + TDW \qquad (1)$$

where TDR is the total data read and TDW is the total data written in bytes.

Attempts have been made to determine how different kinds of workloads affect the survivability of HDDs. However, information pertaining to the amount of random data or sequential data transferred to the disks is not available in this dataset.

Therefore, another parameter, the AAR, is exploited in this investigation. The values of this parameter May vary over time. There could be significant activity at certain times and a lack of activity at other times, depending on application requirements. Keeping in mind the lack of information on the timelines of these cycles in view, a parameter called the average access rate (AAR) is introduced, which is defined as follows:

$$AAR = \frac{TD}{PT} \qquad (2)$$

where, PT is the power-on time in minutes. This parameter intuitively gives a sense of the data transfer rate of a customer application to the HDD. In this analysis, the TD in bytes and the AAR in bytes per minute are used as parameters to measure the survivability of the HDDs.

2 Brief overview of survival analysis

Survival analytic models, which form a branch of statistics, share certain similarities with logistic and linear models. In this analysis, survival analytic models are chosen over logistic or linear models because they consider parameters such as the event time and event probability that are not considered by the alternatives (Kalbfleisch and Prentice, 2011; Lambert and Royston, 2009). The key parameters and associated functions are outlined in this section.

Usually, failed HDDs are returned by different customers at different times. This research takes into account HDDs received over the past 5 years. A detailed analysis was carried out based on the power-on time of the HDDs as well as their workloads. The power-on time is the amount of time for which the HDDs are used in any system. The start time for the analysis on all HDDs is the same, i.e., 0. The end time is the highest value of the power-on time available in the data set. A test suite determines whether one or more events have occurred on an HDD that belongs to a set of failure events. If such an event occurs, the amount of power-on time on that HDD is considered its failure event time.

In survival analysis, the event time distribution is quantified using the following four functions.

2.1 The cumulative distribution function (CDF)

The CDF of hard disk random failure is expressed using a random variable X as

$$F(x) = P(x < X) \qquad (3)$$

where, $F(x)$ is the CDF and the right-hand side represents the probability that X has a value less than or equal to x.

2.2 The probability density function (PDF)

The PDF of a random variable X, denoted $f(x)$, is defined by:

$$f(x) = \frac{d}{dx} F(x) \qquad (4)$$

The PDF is the derivative or slope of the cumulative distribution function (Zhang et al., 2023).

2.3 The survival function

The survival function captures the probability that a component or system is alive and functional beyond a defined point on the time axis (Rincón et al., 2017). Let X be a continuous random variable with cumulative distribution function $F(x)$ in the interval $[0, \infty]$. Its survival function is defined as follows:

$$S(x) = P(\{X > x\}) = 1 - F(x) \qquad (5)$$

The above function helps to define the probability of X being alive just before exceeding duration x, or the probability that the failure event has not occurred at all in the considered x interval. The survival

curve describes the relationship between the probability of survival and time.

2.4 Hazard function

The hazard function h(x) is given by the following equation:

$$h(x) = \frac{f(x)}{S(x)} \tag{6}$$

The hazard function is more intuitive in survival analysis than the probability distribution function because it quantifies the instantaneous risk that an event will take place at time x given that the subject survived to time x [6].

In survival or failure test cases, it is essential to determine whether variables are correlated with survival or failure times. However, this correlation analysis is not simple (Royston and Sauerbrei, 2008; Kalbfleisch and Prentice, 2011; Lambert and Royston, 2009). This is because, most likely, the dependent variable of interest does not follow an exponential distribution but instead a normal distribution. Another contributing factor is incomplete datasets generated from research or analytical studies, where the complete output is not as expected or the dataset is censored.

2.5 The cox proportional hazards model (CPHM)

The CPHM relies on variables that are correlated to survival and does not make presumptions about base line hazard rate of each variable. Therefore, the Cox regression method is much more useful than the Kaplan–Meier estimator (KME) approach, which involves a lot of explanatory variables (Kleinbaum and Klein, 1996; Hosmer et al., 2008; David, 1972). A brief exposition on KME is given below to validate the choice of the CPHM for the proposed research.

In its generic form, the KME can be expressed as below using an example.

Consider a sample size of population N, with t being the time axis. Assume $t1, t2, \ldots ti, \ldots tN$ are the observed lifetimes of the sample size N, such that $t1 \leq t2 \leq t3 \leq \ldots \leq tN$, with $\hat{S}(t)$ the probability that a member has a lifetime exceeding t. In this scenario, the Kaplan–Meier estimator tries to establish a survival function at ti between members who have experienced the event versus members who have not. Let di be the set of members who have experienced the event and ni be the members who are yet to experience the event. Then, the KME can be expressed as:

$$S(t) = \Pi_{t_i \leq t}\left(1 - \frac{d_i}{n_i}\right) \tag{7}$$

Assuming, $\hat{S}(t)$ is the probability that a given member from the sample size has a lifetime exceeding time t.

As no assumptions are made about the nature of the survival distribution, the CPHM model is considered the most generic regression model. Hence, Cox's regression model May be considered

to be a "*semi-parametric*" model (Kalbfleisch and Prentice, 2011). The CPHM is represented as follows:

$$h(t,x) = h_0(t)\exp\left(\Sigma_{i=1}^{p}\beta_i X_i\right) \tag{8}$$

where, $X = (X1, X2, \ldots, Xp)$ are the explanatory/predictor variables, h0(t) is the baseline hazard function, and βi are the regression coefficients.

To linearize this model, both sides of the equation are divided by h0(t) and then the natural logarithm is then taken on both sides. As a result, a fairly "simple" linear model can be readily estimated. The final result is:

$$\log\left(\frac{h(t,x)}{h_0(t)}\right) = \Sigma_{i=1}^{p}\beta_i X_i \tag{9}$$

The CPHM has been endorsed by many researchers as it is quite robust for computing the associated survival probabilities through balancing potential predominant variables (Royston and Sauerbrei, 2008; Kalbfleisch and Prentice, 2011; Lambert and Royston, 2009; Kleinbaum and Klein, 1996; Hosmer et al., 2008). In essence, the CPHM is most suitable for representing and interpreting impact of AAR and TD groups on survival probability of HDD. Also statistically R programming language R is also commensurate tool implementing CPHM for the proposed research (Kleinbaum and Klein, 1996; Hosmer et al., 2008; David, 1972; Kaplan and Meier, 1958; Lane et al., 1986).

3 Modeling and discussions

As the HDDs considered here form part of a storage system, software drives the operation of these hard disks from insertion into the storage system until the drives declare themselves to have failed or software identifies them as faulty (Equations 1–9).

3.1 HDD failure

HDD failure occurs when an HDD is no longer capable of performing IO operations. However, the word HDD failure is somewhat vague, and the threshold for failure is different for each subsystem.

SMART data in its intrinsic form seems to be insufficient because many HDDs undergo unexpected hardware failures, such as head crashes or excessive media errors, without any warning from SMART subsystems. Storage controllers May have different thresholds for declaring an HDD as having failed, such as if a drive gives a certain number of unrecovered errors within a certain period, shows too many timeouts during data access, or cannot reassign a logical block address (LBA).

Therefore, the basis of the present analysis is the event of HDD failure in a storage system. Whenever there is an HDD failure event in a customer data center, the failed HDD is sent to a lab for investigation. The test suites used by storage system vendors mostly rely on the log data available on the HDD, which can be queried through Small

Computer System Interface (SCSI) commands. These test suites have defined thresholds, and if a drive exceeds the thresholds in any certain area, it is declared a failure. There is a never-ending battle between storage controller providers and HDD manufacturers concerning the definition of failure because the manufacturer has to bear the cost if the HDD is still within a warranty period.

Typically, storage systems support multiple HDD models from different manufacturers. Therefore, the gross or initial drive sample considered for this analysis includes different drive models. Out of the various HDD models present in the gross sample, twenty drive models are chosen for this investigation. Historical data collected for these twenty drive models is analyzed and plotted to determine the failure pattern and its dependence on certain parameters. In-depth results are provided for only two drive models, as the other drive models are similar.

3.2 Data description

Data from two hard disk drive models (the M1-A1 model and the M2-A1 model) were analyzed using R. Table 1 represents a sample raw data for model M1-A1(HDD Manufacture's Model Name - NB1000D4450), as given below.

The graphical representation of the output using the power-on time (in minutes) on the x-axis and the failure count on the y-axis were useful for assessing efficacy of the proposed prediction model. To analyze how the workload affects the survivability of an HDD, two attributes are considered: the TD and the AAR. HDDs of the same model were grouped by both TD as well as AAR.

These g number of groups were labeled as TDi where $i = 1, 2, …, g$. The HDDs were also grouped by their average access rates (AAR) into m number of groups called AAR$_j$ where $j = 1, 2, …., m$.

The ranges for these groups were determined automatically in R such that each group had an approximately equal number of drives. A particular drive could belong to any one of the TD groups and any one of the AAR groups, so a drive with low total data could have a high access rate and a drive with high total data could have a low access rate.

To create the total data (TD) accessed till the drive was in operational state and average access rate (AAR) groups for the same period is as in Figure 1.

The flowchart shown above for formation of TD_groups (denoted as *** inside Figure 1) utilizes the **cut2** function in R to divide the variable AAR into groups or intervals. The components of the algorithm are as follows:

- **AAR**: This is the variable contains the AAR data. It represents the ratio of the total data accessed to the power-on time of the hard disk drive.
- **cut2()**: This is a function in R, contained in **Hmisc** package, is used to divide a continuous variable into intervals or groups. Unlike the base R **cut()** function, **cut2()** ensures that each group has approximately the same number of observations.
- **g = n_groups**: This argument specifies the desired number of groups into which the variable AAR is to be divided and n_groups represent the value assigned to the variable indicating the number of groups to be created.
- **levels.mean = TRUE**: This argument indicates that the labels for the intervals or groups should represent the mean value of the observations within each group. This helps with interpreting the intervals more intuitively.
- Overall, the **cut2()** function(denoted as ### inside Figure 1) with the provided arguments divides the AAR variable into n_groups intervals, ensuring that each interval contains approximately the same number of observations, and assigns labels representing the mean value of the observations within each interval.

From the data and plotted graphs (Figures 2A,B), it is observed that there are initially a large number of failures, but later, the rate of failures is constant for some time and then increases again. This behavior seems to correspond to a "bathtub" -type curve.

3.3 Procedure for replicating the model analysis
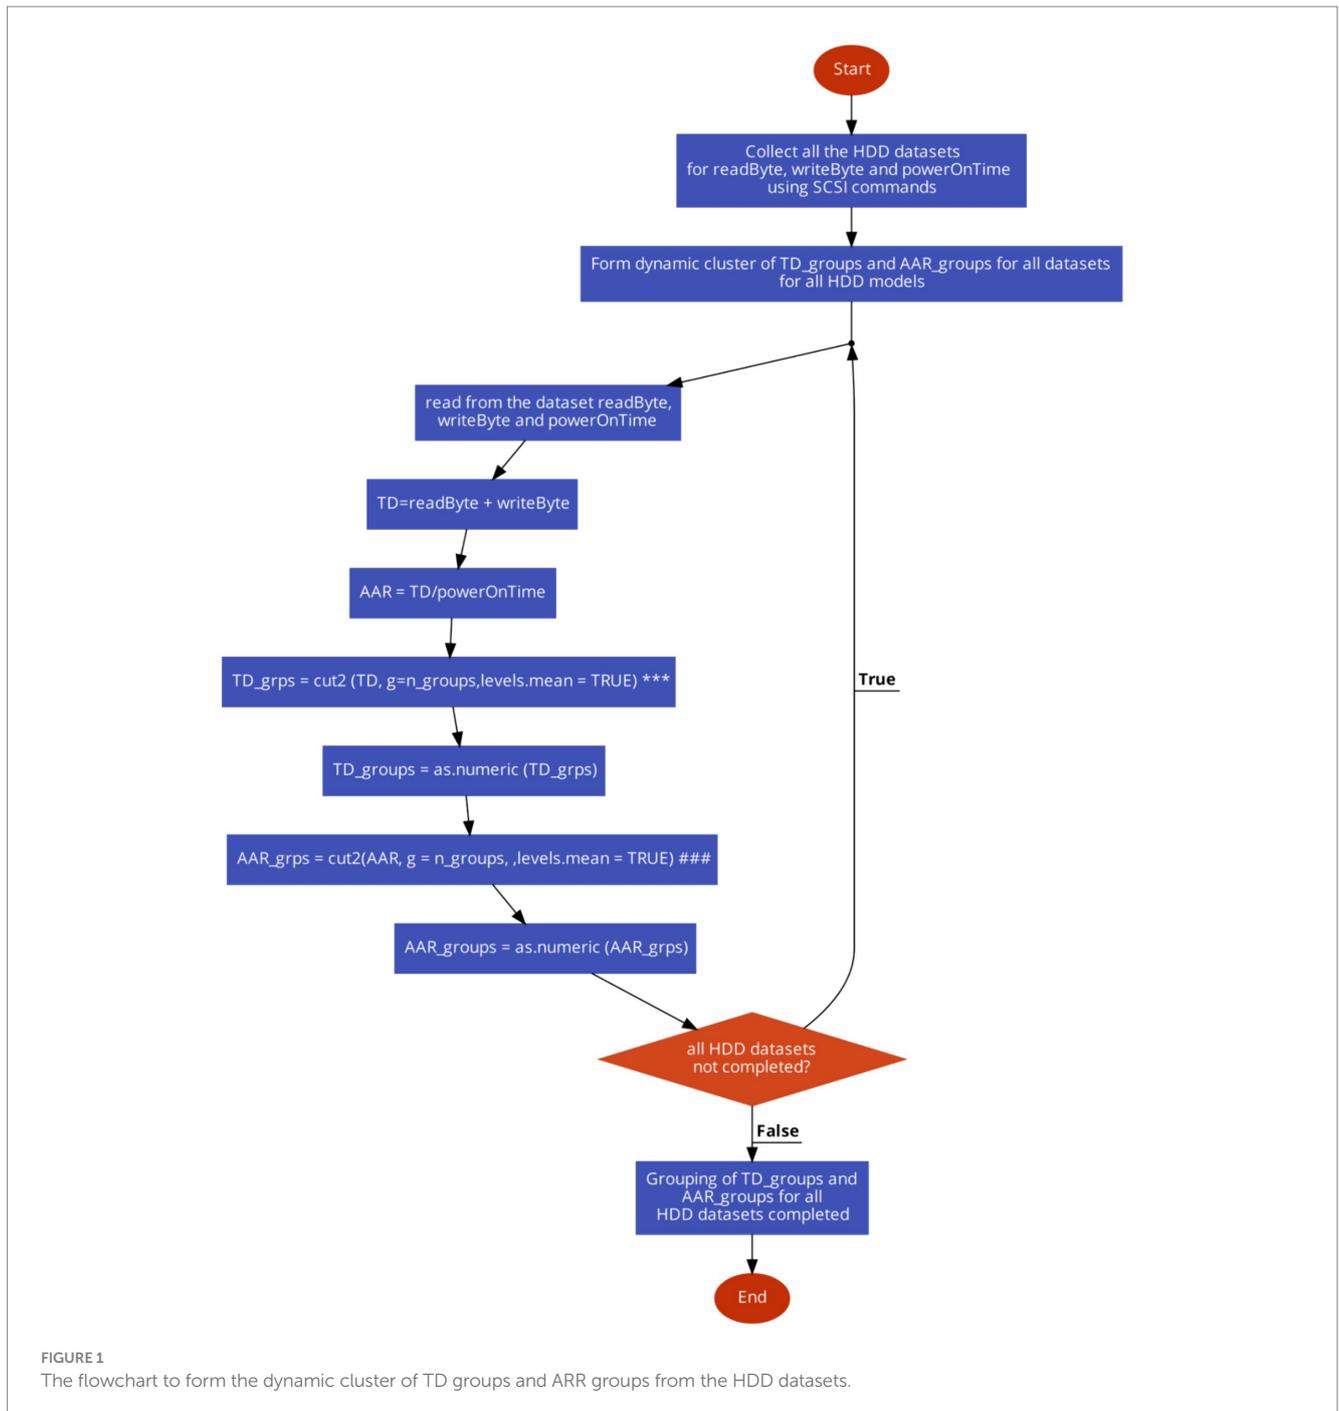
i. Setting Up the Environment.
R is installed on the system along with necessary libraries.
ii. Reading and Formatting Raw Data.
The CSV file containing the dataset is read ensuring appropriate file path.

TABLE 1 Sample raw data (The highlighted yellow columns represent the data used in the current analysis) of HDD model.

Name	HDD Model	FW	Write Rec Err	WriteBytes	Read Rec Err	ReadBytes	Power on time	SSUITE status
6001r.txt	NB1000D4450	XX04	0×0000	0x0f60381db400	0x685eba9c	0x2e739518ec00	640,975	0
6002r.txt	NB1000D4450	XX06	0×0000	0x00ced22e8800	0x36257e	0x0061ba630800	10,030	0
6005r.txt	NB1000D4450	XX06	0×0000	0x04d085cf5200	0x82bbf6	0x004f4faa9c00	626,385	0
6007r.txt	NB1000D4450	XX06	0×0000	0×000032800	0×0035	0x0000203e00	11	1
6007r.txt	NB1000D4450	XX06	0×0000	0x01bd86913c00	0×2284569	0x00682858ca00	403,928	0
6009r.txt	NB1000D4450	XX06	0×0000	0x02b9bb3ab000	0x171ce42f	0x11ce58d46400	225,990	0
6010r.txt	NB1000D4450	XX06	0×0000	0x05415ac38800	0x3e9a2c69	0x0930f8101800	233,520	0
6010r.txt	NB1000D4450	XX05	0×0000	0x03b3d4f68e00	0x96ba9570	0x0f8ca67b10	677,146	0
6012r.txt	NB1000D4450	XX06	0×0000	0×000029000	0×0007	0x000010bc00	4	1
6013r.txt	NB1000D4450	XX04	0×0000	0x12a77d1a0c00	0x95e8008	0x2765ae202e00	359,596	0

FIGURE 1
The flowchart to form the dynamic cluster of TD groups and ARR groups from the HDD datasets.

The SSUITEstatus column is formatted in order to convert it to numeric.

iii. Grouping Data Based on Total Data and Data Rate.

Total Data (TD) and Average Access Rate (AAR) are calculated. TD and AAR are grouped into specified number of groups.

iv. Converting Data frame to Data Table.

The dataframe is converted to a data.table for efficient data manipulation:

v. Calculating Cumulative Sum and Percentage Failures.

Cumulative sums and percentage failures by AAR_groups are calculated.

vi. Splitting Data into Passed and Failed Subsets.

Subsets of the data for passed and failed tests are created. Cumulative sums and percentages for both subsets are calculated.

vii. Creation of a Survival Object.

A survival object using PowerOnTime and SSUITEstatus is created.

viii. Creation of Cox Proportional Hazards Model.

The Cox model is created using AAR_groups and TD_groups.

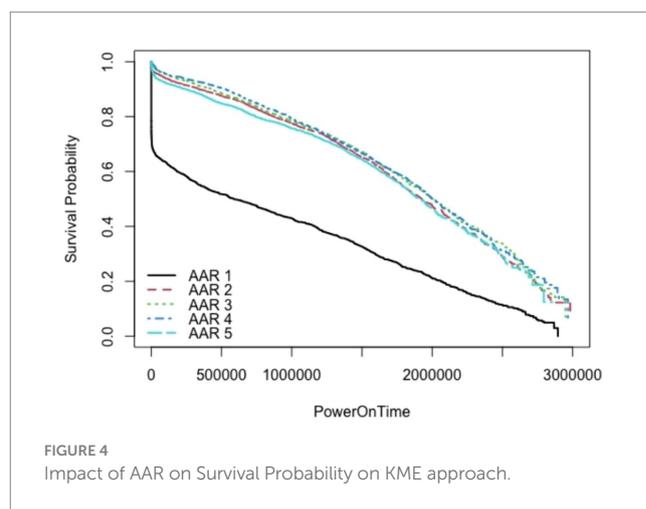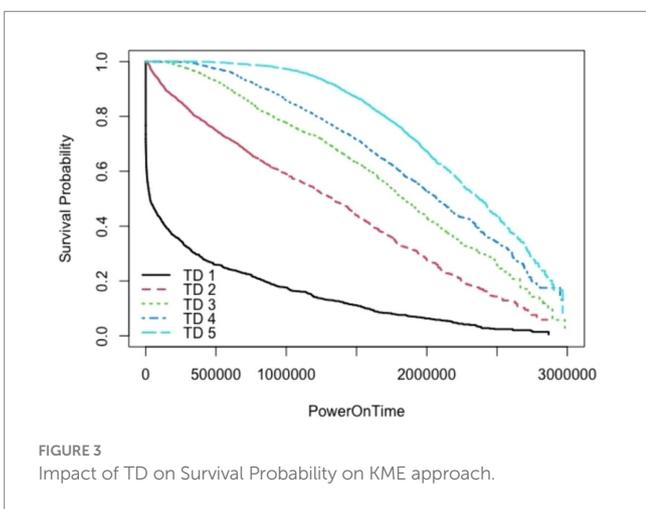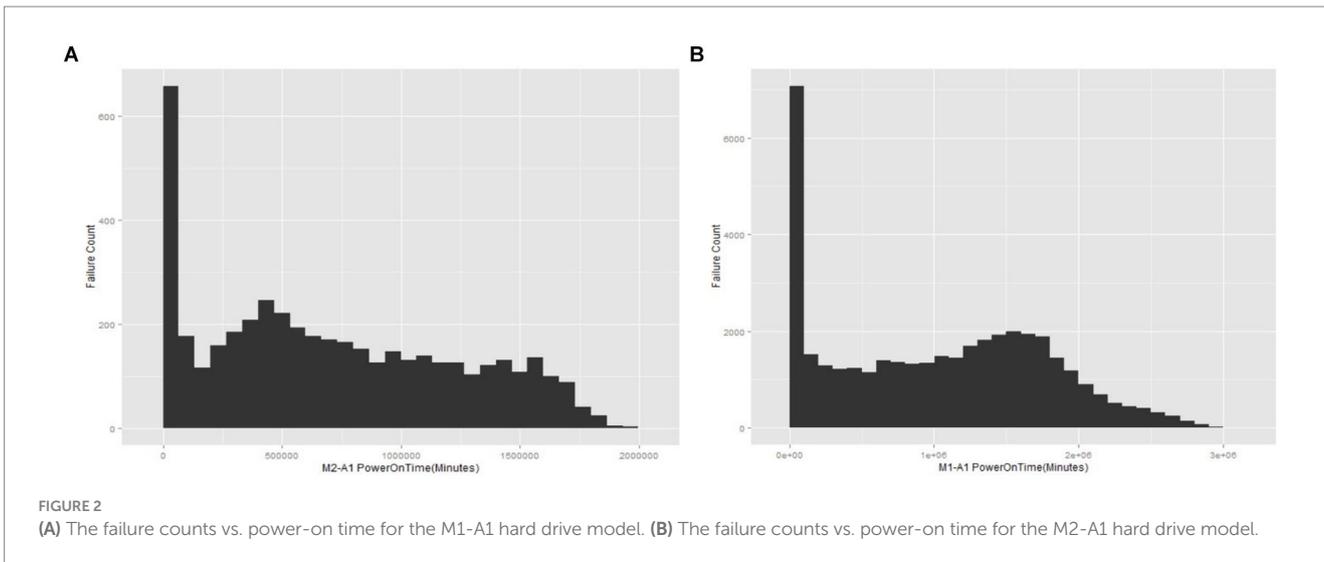ix. Checking of Proportional Hazards Assumption.

cox.zph is used to check the proportional hazards assumption.

x. Predictions Based on the created Cox Model.

The output is predicted based given inputs using Cox model.

xi. Display of the Output.

The predicted pass or fail value is printed.

FIGURE 2
(A) The failure counts vs. power-on time for the M1-A1 hard drive model. **(B)** The failure counts vs. power-on time for the M2-A1 hard drive model.

FIGURE 3
Impact of TD on Survival Probability on KME approach.

FIGURE 4
Impact of AAR on Survival Probability on KME approach.

4 Case studies and comparison of results

Previous research examined the susceptibility to failure of the new generation of hard disks because of a reduction in the spacing between the heads and media. The research also examined the effect of the amount of data written to or read from the drive. The current research attempts to evaluate the effect of the data access rate on the survivability of the HDDs using the Cox model.
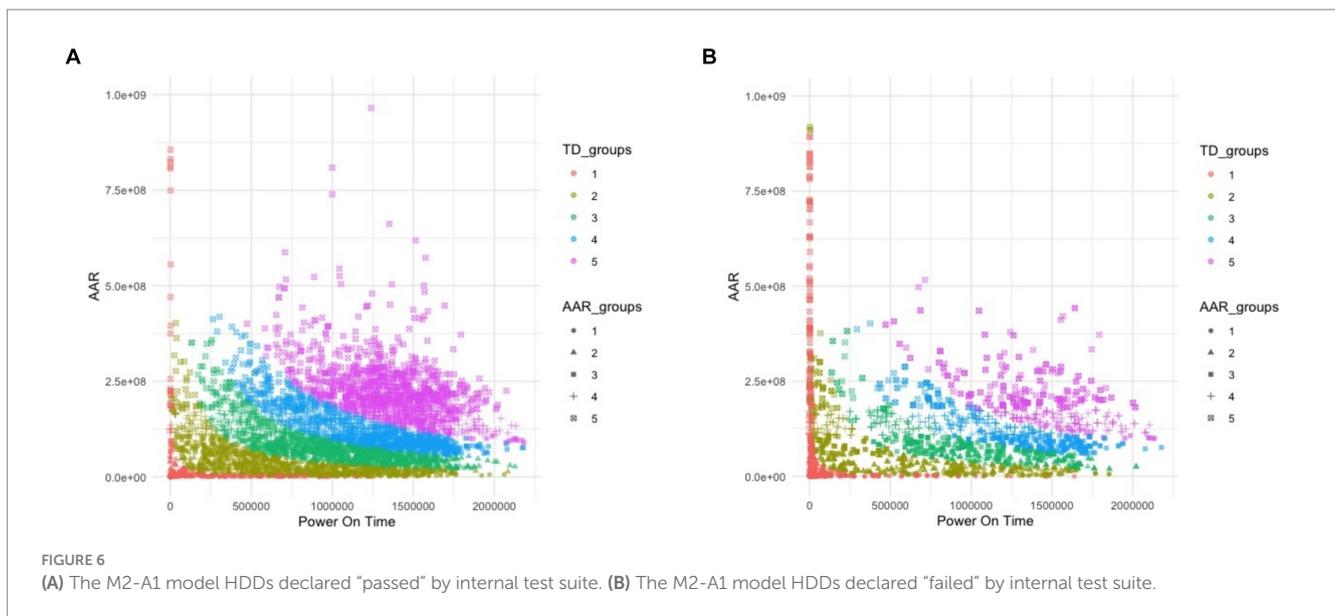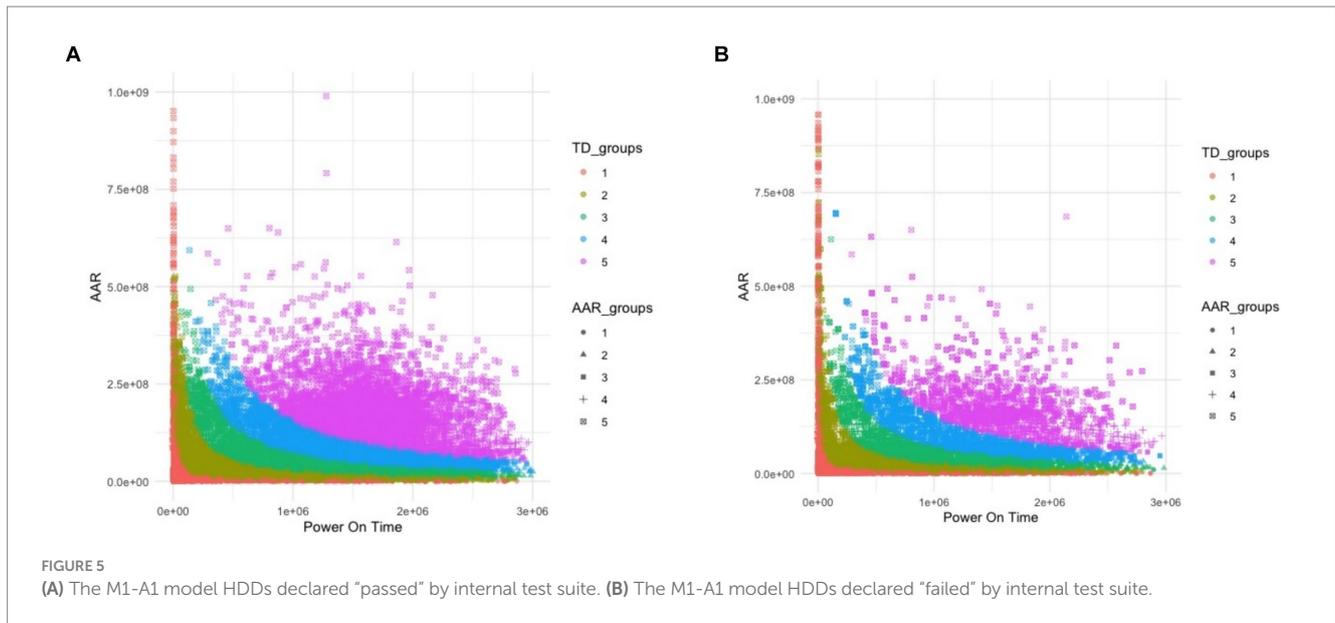
4.1 SMART data based model for internal test suit

SMART is a standard technology embedded in most modern hard disk drives (HDDs) to monitor various indicators of drive health. These indicators include attributes, like read error rates, spin-up time, temperature, and reallocated sectors, count. The primary goal of SMART is to predict drive failures before they occur, allowing for preventive measures such as data backup or drive replacement. The SMART system has been reported in literature (Villalobos, 2020;

Rajashekarappa and Sunjiv Soyjaudah, 2011) for HDD failure prediction. The internal analysis for the SMART based data has employed Kaplan-Meir approach. The impact of TD and AAR are exhibited in the following Figures 3, 4.

SMART is widely used and supported across the industry, making it a reliable standard for monitoring HDD health (Villalobos, 2020; Rajashekarappa and Sunjiv Soyjaudah, 2011). It provides early warnings based on predefined thresholds for various health indicators, enabling proactive maintenance. However, it is limited by the predefined thresholds and May not account for all failure modes; sometimes produces false positives or misses failures. It is evident from the above figures that KME is well suited for incorporating one impact of TD on survival probability. The impact of AAR is also reflected in survivability in Figure 4. However, the simultaneous two impact of both TD and ARR is not possible to capture using KME approach.

Figures 5A,B, 6A,B present the data for HDDs with model parameters of interest in this analysis, namely, TD and AAR. The color code indicates which groups the hard drives belong to and the total amount of data read by or written to them. Figures 5, 6 show AAR and TD distributions of the drive population. The dispersion of the data

FIGURE 5
(A) The M1-A1 model HDDs declared "passed" by internal test suite. **(B)** The M1-A1 model HDDs declared "failed" by internal test suite.

FIGURE 6
(A) The M2-A1 model HDDs declared "passed" by internal test suite. **(B)** The M2-A1 model HDDs declared "failed" by internal test suite.

used for the M1-A1 model is given below in Table 2 (for AAR_groups) and Table 3 (for TD_groups).

Some interesting inferences can be made from the plots shown in Figures 5, 6. Figure 5A shows the data for M1-A1 HDDs that failed in the customer environment as soon as a certain rate of access was reached but passed the internal test suite run in the lab. The drives were declared failed either by the storage controllers that were using them or by the test suites run on them later. This shows that each drive has an access rate threshold point, after which the probability of failure increases drastically.

Figure 5B presents the data for M1-A1 HDDs that failed in the customer environment and failed in the internal suite. Figure 6A shows the data for M2-A1 HDDs that failed in the customer environment as soon as a certain rate of access was reached but passed the internal test suite run in the lab. Figure 6B shows the data for the M2-A1 HDDs that failed in the customer environment and failed in the internal suite. These figures show that the failures are

similar between groups with different amounts of data. After a certain rate of access, the majority of the drives failed regardless of the amount of data accessed.

4.2 Survivability analysis using cox proportionality model

The proposed model is a predictive model based on Total Data (TD) read/written and Average Access Rate (AAR). It introduces a novel strategy of reallocating HDDs with high failure probabilities to different redundancy groups.

It provides with a proactive approach to HDD management by predicting failures based on data access patterns. It enhances system resilience by dynamically reallocating resources based on failure probabilities. Also, it utilizes real-world data from storage controllers for validation, ensuring practical applicability.

TABLE 2 Dispersion of AAR_groups the dataset for M1-A1 used by the internal test suite.

Data Sets	Count	AAR groups	Mean	Median	Minimum	Maximum	Standard deviation
Total dataset	7,903	1	3,102,686	469603.1	5.35E-02	1.24E+07	3.92E+06
	7,903	2	26,100,943	25884602.9	1.24E+07	4.04E+07	7.99E+06
	7,903	3	58,562,114	57,867,575	4.04E+07	7.90E+07	1.12E+07
	7,903	4	104,903,206	103477796.4	7.90E+07	1.36E+08	1.67E+07
	7,902	5	1.4613E+10	188,703,828	1.36E+08	9.34E+13	1.07E+12
Dataset declared as Failed by internal test suite	4,864	1	2,402,446	228,134	5.35E-02	1.24E+07	3.59E+06
	3,020	2	25,880,789	25,160,218	1.24E+07	4.04E+07	7.99E+06
	2,932	3	58,695,593	58,137,170	4.04E+07	7.90E+07	1.13E+07
	2,761	4	104,873,553	103,724,219	7.90E+07	1.36E+08	1.66E+07
	2,862	5	3.4415E+10	189,678,100	1.36E+08	9.34E+13	1.75E+12
Dataset declared as passed by Internal test suite	3,039	1	4,223,439	2,980,341	5.66E-02	1.24E+07	4,155,452
	4,883	2	26,237,101	26,245,474	1.24E+07	4.04E+07	7,985,086
	4,971	3	58,483,386	57,727,494	4.04E+07	7.90E+07	11,098,270
	5,142	4	104,919,129	103,317,505	7.90E+07	1.36E+08	16,738,192
	5,040	5	3,368,262,156	188,156,025	1.36E+08	1.56E+13	2.1924E+11

TABLE 3 Dispersion of TD_groups the dataset for M1-A1 used by the internal test suite.

Data Sets	Counts	TD groups	Mean	Median	Minimum	Maximum	Standard Deviation
Total dataset	7,903	1	8.80E+13	4.67E+13	123,392	1.97E+15	1.12E+14
	7,903	2	8.80E+13	4.67E+13	123,392	1.97E+15	1.12E+14
	7,903	3	8.80E+13	4.67E+13	123,392	1.97E+15	1.12E+14
	7,903	4	8.80E+13	4.67E+13	123,392	1.97E+15	1.12E+14
	7,902	5	8.80E+13	4.67E+13	123,392	1.97E+15	1.12E+14
Dataset declared as failed by internal test suite	5,211	1	8.80E+13	4.67E+13	1.23E+05	1.97E+15	1.12E+14
	3,337	2	8.80E+13	4.67E+13	1.23E+05	1.97E+15	1.12E+14
	2,927	3	8.80E+13	4.67E+13	1.23E+05	1.97E+15	1.12E+14
	2,777	4	8.80E+13	4.67E+13	1.23E+05	1.97E+15	1.12E+14
	2,187	5	8.80E+13	4.67E+13	123,392	1.97E+15	1.12E+14
Dataset declared as passed by internal test suite	2,692	1	8.80E+13	4.67E+13	1.23E+05	1.97E+15	1.12E+14
	4,566	2	8.80E+13	4.67E+13	1.23E+05	1.97E+15	1.12E+14
	4,976	3	8.80E+13	4.67E+13	1.23E+05	1.97E+15	1.12E+14
	5,126	4	8.80E+13	4.67E+13	1.23E+05	1.97E+15	1.12E+14
	5,715	5	8.80E+13	4.67E+13	1.23E+05	1.97E+15	1.12E+14

However, it May need customization for different storage environments and workloads.

The Cox Model uses the parameters of the AAR groups and TD groups to generate the survival model. The equation for the Cox model using R is (Equation 10):

$$M1A1\_cox = coxph(M1A1\_surv \sim as.\ factor(AAR\_groups) + as.\ factor(TD\_groups),\ data = M1\_A1) \qquad (10)$$

The Cox model equation fits a Cox proportional hazards regression model (**coxph()**) to the survival data (**M1A1_surv**). It includes the average access rate (**AAR_groups**) and the total data groups (**TD_groups**) as covariates. The model aims to understand how these factors influence the hazard rate or survival probability over time. The output, **M1A1_cox**, is a Cox model object containing the results of the regression analysis.

The effect of the AAR on each of these TD groups was analyzed. For a given AAR, the survival data were generated for each of the TD groups using the above Cox model, whose algorithm is outlined in Figure 7.

For each value of AAR, survival data were generated for the different data groups, using the algorithm given below in Figure 7.

The algorithm in Figure 7, as demonstrated using ARR1 values for different TD groups (TD1-TD5), was applied in a similar way to generate the survival graphs for the other AAR groups, AAR2–AAR5. It utilizes cut2() for formation of TD_groups (denoted as *** inside Figure 7) and also for AAR_groups (denoted as ### inside Figure 7). In the algorithms, the HDD model M1-A1 was used as an example.
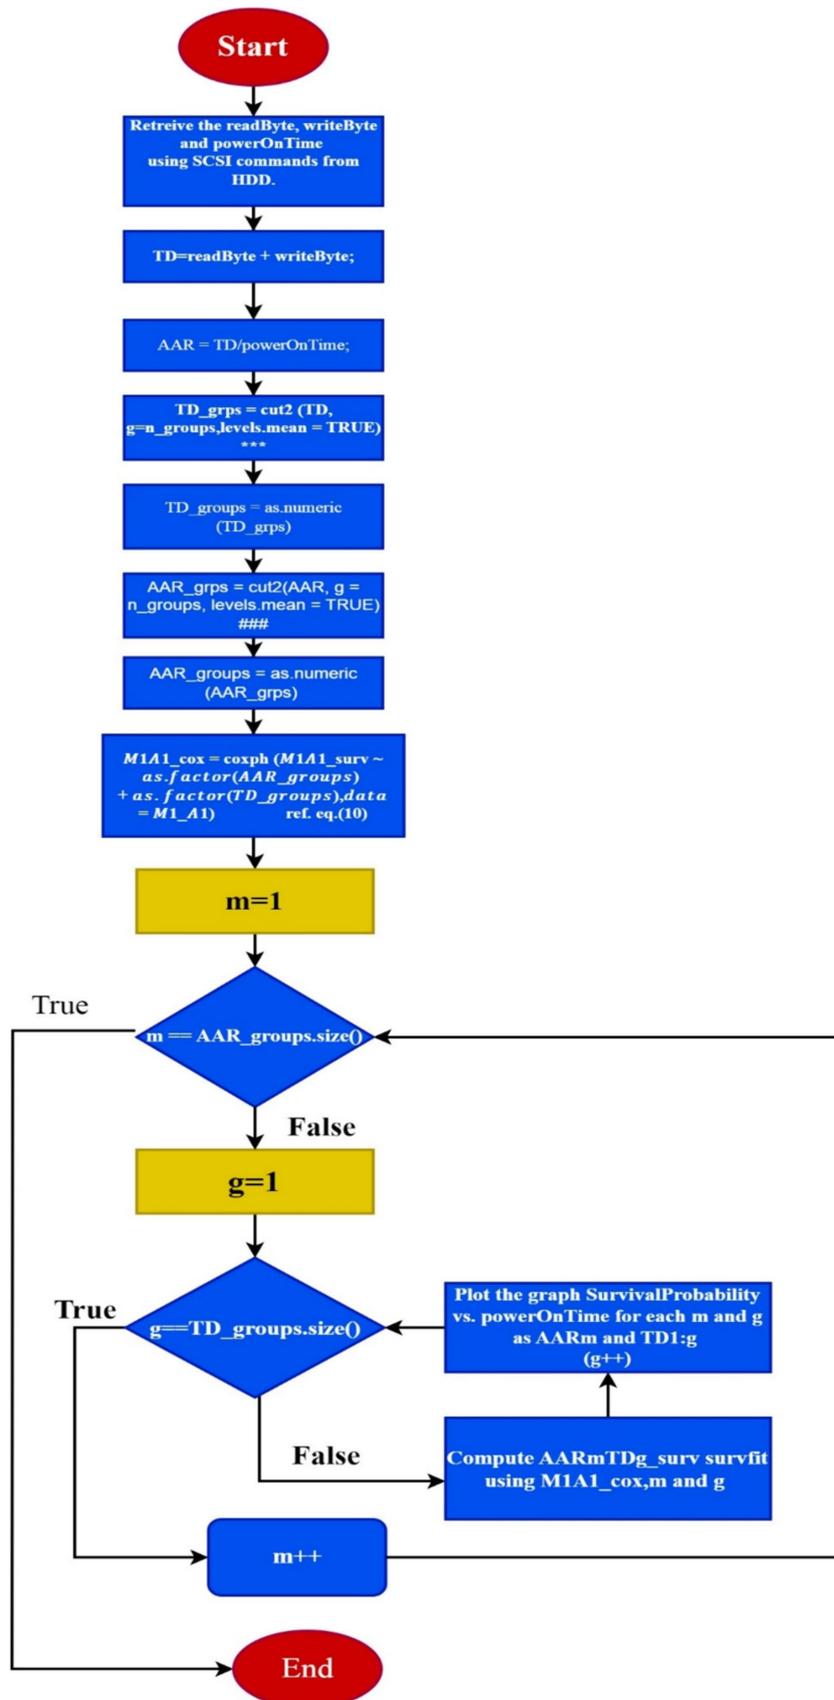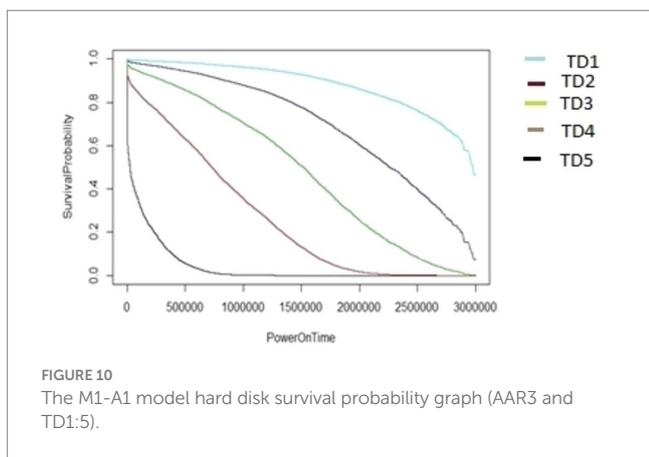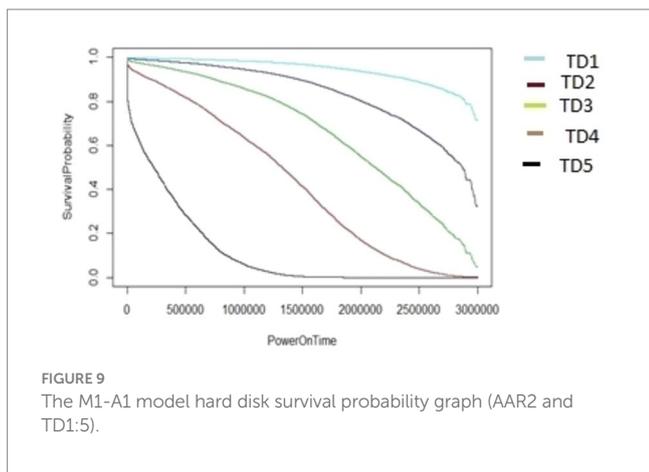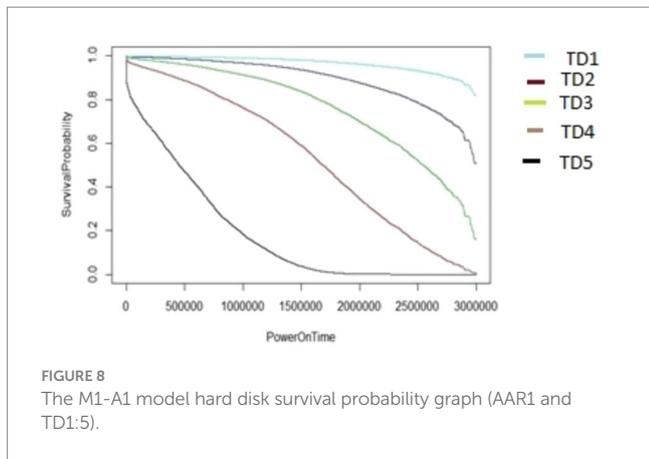
FIGURE 7
The flowchart for the validation of the survivability-based predictive data model.

The same algorithms can be used for M2-A2 model drives to generate survival data and survival graphs. The unit of the PowerOnTime (X-axis) in Figures 8–17 is in minutes. The effect of workload on the survivability of a hard disk (M1-A1 model).

FIGURE 8
The M1-A1 model hard disk survival probability graph (AAR1 and TD1:5).

FIGURE 9
The M1-A1 model hard disk survival probability graph (AAR2 and TD1:5).

FIGURE 10
The M1-A1 model hard disk survival probability graph (AAR3 and TD1:5).

FIGURE 11
The M1-A1 model hard disk survival probability graph (AAR4 and TD1:5).

FIGURE 12
The M1-A1 model hard disk survival probability graph (AAR5 and TD1:5).

FIGURE 13
The M2-A1 model hard disk survival probability graph (AAR1 and TD1:5).
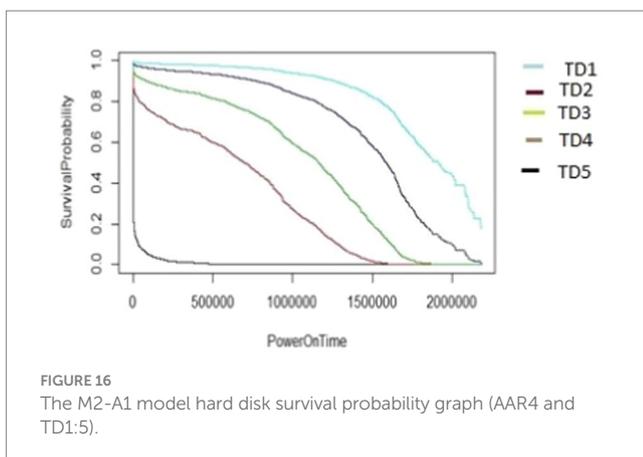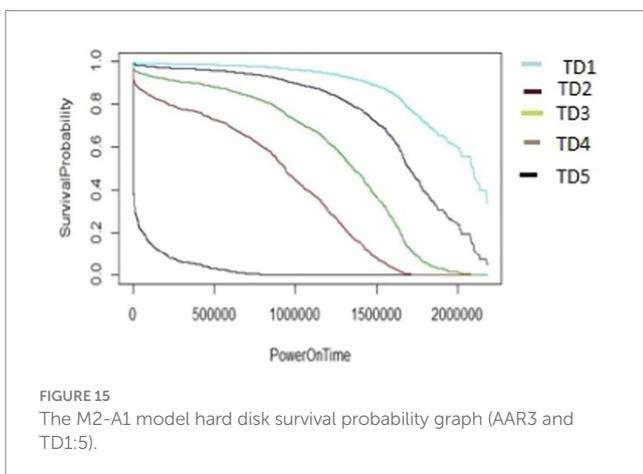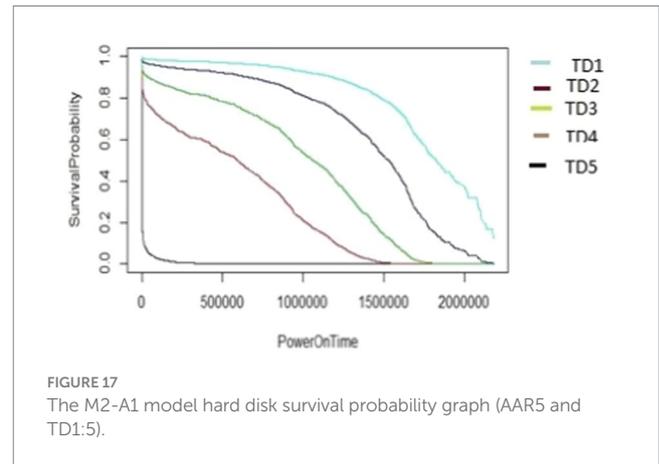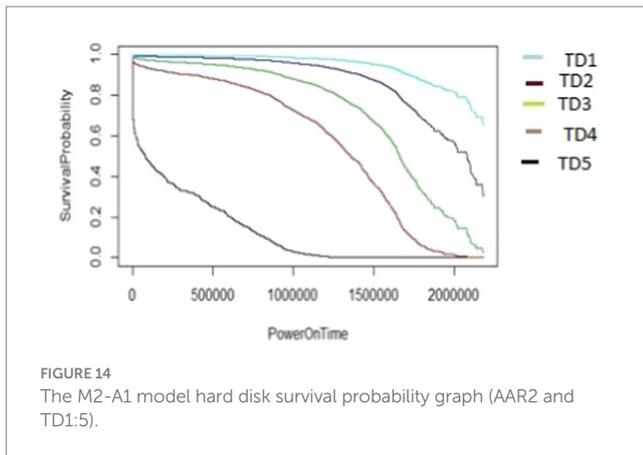
In summary, the proposed method stands out by integrating workload-specific metrics (TD and AAR) into the predictive model, which offers a more targeted approach to identifying potential HDD
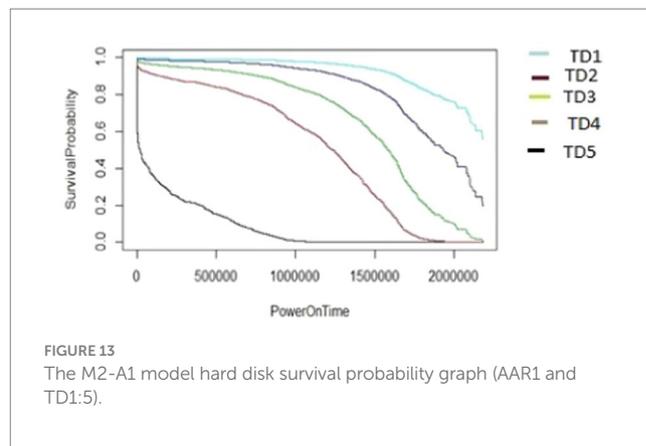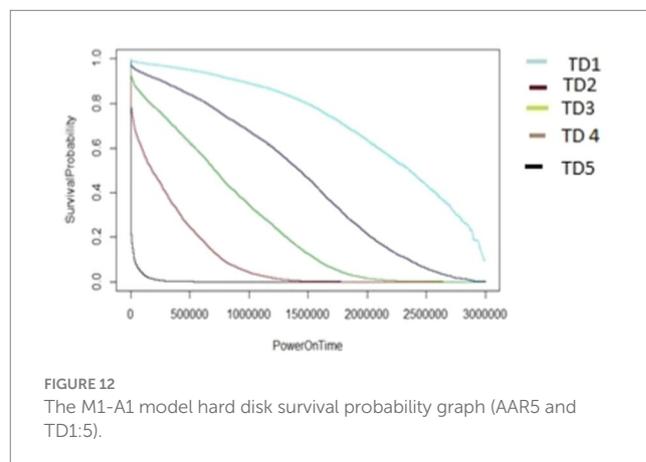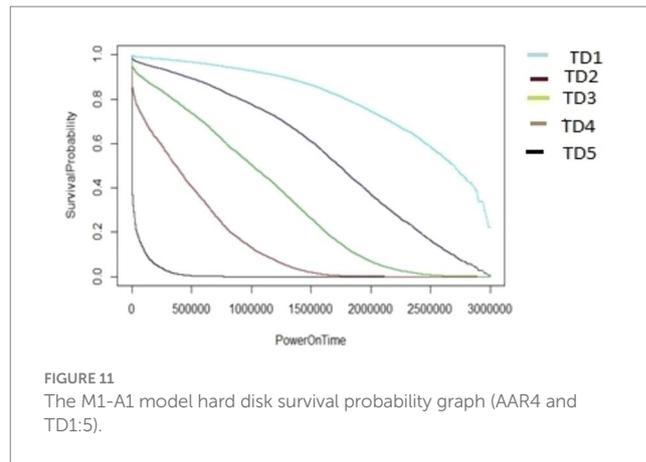
failures compared to traditional SMART models. By focusing on empirical data from actual storage systems and introducing dynamic reallocation strategies, this method aims to provide a more practical and effective solution for mitigating HDD failures in enterprise environments.

This comparison highlights the advantages of the proposed method in enhancing system resilience and optimizing resource allocation, while also acknowledging the challenges and areas for further improvement. The authors emphasize that this comprehensive approach can significantly contribute to the advancement of storage system reliability and efficiency.

FIGURE 14
The M2-A1 model hard disk survival probability graph (AAR2 and TD1:5).

FIGURE 17
The M2-A1 model hard disk survival probability graph (AAR5 and TD1:5).

FIGURE 15
The M2-A1 model hard disk survival probability graph (AAR3 and TD1:5).

FIGURE 16
The M2-A1 model hard disk survival probability graph (AAR4 and TD1:5).

The graphs shown in Figures 8–12 were plotted using data for the HDD model M1-A1. They show that for a given data volume (TD), as the access rate is changed from AAR1 to AAR5, the survival probabilities are reduced. They also show that the access rate has a significant effect on HDD failures.

The graphs shown in Figures 8–12 also confirm that for a fixed amount of data, if the data access rate is increased, the survival probability of the drives decreases. They also show that the rate of access has a significant impact on survivability apart from just the total data written or read.

The graphs shown in Figures 13–17 were plotted using data from the HDD model M2-A1. They show that for a given amount of data (TD), as the access rate changes from AAR1 to AAR5, the survival probabilities are reduced. They also show that the access rate has a significant effect on HDD failures. The graphs shown in Figures 12–16 confirm that, for a fixed amount of data, if the data access rate is increased, the survival probability of the drives decreases. They also show that the rate of access has a significant impact on survivability apart from just the total data written or read.

5 Conclusion

In the modern digital world, data is one of the critical assets of any business, and data availability and data access are important metrics for predicting the failure of HDDs. The method for storage system controller firmware proposed in this work could be used to manage failure-predicted disk drives efficiently and intelligently, in order to provide greater data availability and survivability to customers. This approach allows one to apply better throttling mechanisms and the preemptive migration of data from HDDs that are predicted to fail. The proposed model paves the way for facilitating collaborations between storage experts, workload analysts, and system architects to merge domain-specific insights. This interdisciplinary approach could foster innovation by integrating expertise from various fields, leading to more robust and comprehensive survivability analyses. Thus, this work should make significant contributions toward improving data availability in companies.

Data availability statement

The datasets presented in this article are not readily available because these datasets are vendor specific. Requests to access the datasets should be directed to amanmallik01@gmail.com.

Author contributions

Funding

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

David, C. R. (1972). Regression models and life tables (with discussion). *J. R. Stat. Soc.* 34, 187–202. doi: 10.1111/j.2517-6161.1972.tb00899.x

George, T. (2013). Why specify workload?. *Western Digital technologies, Inc.,* 3355 Michelson drive. Suite 100 Irvine. California.

Hamerly, G., and Elkan, C. (2001). Bayesian approaches to failure prediction for disk drives. *ICML* 1, 202–209.

Hosmer, D. W., Lemeshow, S., and May, S. (2008). Applied survival analysis: Regression modeling of time-to-event data. Hoboken, New Jersey, USA: John Wiley & Sons.

Kalbfleisch, J. D., and Prentice, R. L. (2011). The statistical analysis of failure time data. Hoboken, New Jersey, USA: John Wiley & Sons.

Kaplan, E. L., and Meier, P. (1958). Nonparametric estimation from incomplete observations. *J. Am. Stat. Assoc.* 53, 457–481. doi: 10.1080/01621459.1958.10501452

Kleinbaum, D. G., and Klein, M. (1996). Survival analysis a self-learning text. *Springer* 52:1528. doi: 10.2307/2532873

Lambert, P. C., and Royston, P. (2009). Further development of flexible parametric models for survival analysis. *Stata J.* 9, 265–290. doi: 10.1177/1536867X0900900206

Lane, W. R., Looney, S. W., and Wansley, J. W. (1986). An application of the cox proportional hazards model to bank failure. *J. Bank. Financ.* 10, 511–531. doi: 10.1016/S0378-4266(86)80003-6

Liu, Q., and Xing, L. (2020). Survivability and vulnerability analysis of cloud RAID systems under disk faults and attacks. *Int. J. Math. Eng. Manag. Sci.* 6, 15–29. doi: 10.33889/ijmems.2021.6.1.003

Mohanta, T., and Ananthamurthy, S. (2006). A method to predict hard disk failures using SMART monitored parameters. Recent developments in National Seminar on devices, circuits and communication (NASDEC2-06). Ranchi. India 243–246.

Mohanta, T., Basireddy, R. R., Nayaka, B. G., and Chirumamilla, N. (2015). Method for higher availability of the SSD storage system. In 2015 international conference on computational intelligence and networks (pp. 130–135). IEEE.

Murray, J. F., Hughes, G. F., and Kreutz-Delgado, K. (eds.) (2003). Hard drive failure prediction using non-parametric statistical methods.

Rajashekarappa, S, and Sunjiv Soyjaudah, K. M. (2011). Self monitoring analysis and reporting technology (SMART) Copyback. In International conference on information processing (pp. 463–469). Berlin, Heidelberg: springer Berlin Heidelberg.

Rincón, C. A., Pâris, J. F., Vilalta, R., Cheng, A. M., and Long, D. D. (2017) Disk failure prediction in heterogeneous environments. In 2017 international symposium on performance evaluation of computer and telecommunication systems (SPECTS) (pp. 1–7). IEEE.

Royston, P., and Sauerbrei, W. (2008). Multivariable model-building: A pragmatic approach to regression anaylsis based on fractional polynomials for modelling continuous variables. Hoboken, New Jersey, USA: John Wiley & Sons.

Shen, J., Wan, J., Lim, S. J., and Yu, L. (2018). Random-forest-based failure prediction for hard disk drives. *Int. J. Distrib. Sensor Netw.* 14:155014771880648. doi: 10.1177/1550147718806480

Smith, T., and Smith, B. (2001). Survival analysis and the application of Cox's proportional hazards modeling using SAS. SAS conference proceedings:. SAS Users Group International.

Smith, T., and Smith, B., (2004).Kaplan Meier and cox proportional hazards modeling: hands on survival analysis. SAS conference proceedings: Western users of SAS software 2004 Pasadena, California.

Smith, T., and Smith, B. (2005). Graphing the probability of event as a function of time using survivor function estimates and the SAS® System's PROC PHREG. In SAS Conference Proceedings: Western Users of SAS Software (pp. 21–23).

Villalobos, C. C., (2020) Survival analysis: hard drive reliability sample. Available at: https://github.com/cconejov/Surv_Analysis_HDD (Accessed November 11, 2023).

Zhang, M., Ge, W., Tang, R., and Liu, P. (2023). Hard disk failure prediction based on blending ensemble learning. *Appl. Sci.* 13:3288. doi: 10.3390/app13053288