Deep Learning-Based Real-Time Failure Detection of Storage Devices

Chuan-Jun Su(✉), Lien-Chung Tsai, Shi-Feng Huang, and Yi Li

135 Yuan-Tung Road, Chung-Li District, Taoyuan City 32003, Taiwan, R.O.C.

Abstract. With the rapid development of cloud technologies, evaluating cloud-based services has emerged as a critical consideration for data center storage system reliability, and ensuring such reliability is the primary priority for such centers. Therefore, a mechanism by which data centers can automatically monitor and perform predictive maintenance to prevent hard disk failures can effectively improve the reliability of cloud services. This study develops an alarm system for self-monitoring hard drives that provides fault prediction for hard disk failure. Combined with big data analysis and deep learning technologies, machine fault pre-diagnosis technology is used as the starting point for fault warning. Finally, a predictive model is constructed using Long and Short Term Memory (LSTM) Neural Networks for Recurrent Neural Networks (RNN). The resulting monitoring process provides condition monitoring and fault diagnosis for equipment which can diagnose abnormalities before failure, thus ensuring optimal equipment operation.

Keywords: Big data · Hard disk · Failure prediction · Recurrent neural networks (RNN) · Long and short term memory (LSTM)

1 Introduction

Following the rapid development of cloud computing technologies, many companies now provide large-scale cloud services to user bases numbering in the millions. Thus, from the service provider's perspective, data center reliability is a critical issue, particularly in terms of hard disk failure on servers, which can interrupt cloud services and result in data loss.

Developers and IT teams require reliable and robust system equipment, including hard drives. Hard drive failure and subsequent data loss are still a serious issue, despite the development of solid state drives (SSD) for servers. SSDs are still considerably more expensive than hard disk drives (HDDs) in server environments. Current HDD providers offer warranty periods ranging from one to three years, but statistics show that as many as 9% of HDDs fail annually [1] without warning. Extended usage of an HDD can result in damage that can result in abrupt failure, potentially leading to data loss with serious consequences for the enterprise. Therefore, maintaining such devices and preventing such loss is of paramount importance.

Traditionally, this is accomplished through preventative maintenance with planned, regular maintenance scheduled based on the age of the equipment. However, this approach can increase operating costs through the pre-mature replacement of still-

T. Ahram (Ed.): AHFE 2019, AISC 965, pp. 160–168, 2020.
https://doi.org/10.1007/978-3-030-20454-9_16

functioning equipment. Predictive maintenance uses actual real-time data to predict future failure of still operating equipment, allowing maintenance personnel to schedule maintenance in advance. Predictive maintenance models are trained using historical data, providing for real-time insight into actual equipment status and predicting future failure. This allows for equipment to be replaced closer to its actual failure point, thus conserving resources.

This study uses big data analytics and artificial intelligence to develop a real-time system for predicting hard drive failure, thus aiding developers and IT teams to maintain data centers. The system can alert maintenance staff to potential imminent failure, allowing for early monitoring and maintenance work. We develop predictive frameworks for hard disk failures by analyzing machine log files instead of using traditional statistical prediction methods. The development process involves the following tasks: (1) retrieving and pre-processing data from the hard disk; (2) training and verifying the predictive model; (3) establishing a predictive system based on the model; and (4) predicting hard drive faults in real-time through streaming data.

The remainder of this paper is organized as follows. Section 2 reviews the relevant literature on techniques such as hard disk failure prediction, recurrent neural networks, and long short-term memory. Section 3 introduces the experimental design, with experimental methods designed based on the requirements and methods introduced in Sects. 1 and 2. Finally, Sect. 4 provides conclusions and recommendations for future research.

2 Related Works

2.1 Predictive Maintenance

Predictive maintenance is one way to maintain industrial, commercial, government and residential technology installations. It involves performing functional inspections and repairs or replacement of components, equipment or devices to maintain machine or facility operations. Maintenance techniques can include restorative, preventive or predictive maintenance.

Predictive maintenance monitors the actual operating condition of equipment and seeking to predict whether it will fail in the future, rather than operating-time or calendar-based maintenance scheduling. PM uses the analysis and modeling of collected equipment data for vibration, temperature and other parameters to predict equipment failure and to devise appropriate maintenance planning, thus reducing maintenance costs. Moreover, continuously collected data can be subjected to big data analysis techniques to incrementally improve fault prediction accuracy.

Canizo [2] established a cloud-based analytics platform that uses data from cloud deployments to generate a predictive model for individual monitored wind turbines, predicting wind turbine status in ten minute increments and displaying a visualization of the predicted results. Zhao [3] proposed a predictive maintenance method based on anomalies detected through data correlations among sensors. The results show this approach outperforms the use of sensor data alone, and can be used to predict failures in advance, thus reducing downtime and maintenance costs.

2.2 Machine Learning

Machine learning is a discipline in computer science that studies the development of programs that can learn. According to Samuel [4], "Machine learning allows a computer to learn without explicit programming." Machine learning techniques are widely used in data mining, email filtering, computer vision, natural language processing (NLP), optical character recognition (OCR), biometrics, search engines, medical diagnostics, credit card fraud detection and speech recognition.

Machine learning is part of artificial intelligence, and is often applied to research on prediction or categorization problems. Machine learning is primarily an algorithm that allows computers to "learn" autonomously. Applying machine learning techniques to historical data allows the computer to learn from feedback, examples and expertise. This feedback loop is used to arrive at identical or completely different predictions for similar future situations. Sample data is used in machine learning training to generate a model, which can then be applied to a test data set for prediction.

The performance limitations of machine learning are determined by data and features, and this upper limit can only be approximated by modeling methods. Prior to machine learning modeling, one must first use feature engineering methods to identify the important eigenvalues in the data. Brownlee [5] argued that feature engineering is the process of transforming raw data into features that better represent the underlying problems of predictive models, thereby improving model accuracy for invisible data. However, using machine learning methods to predict possible hard disk failure requires expert knowledge to perform feature engineering. With the important eigenvalues and data, the prediction module can be built, but hard disk data will vary between manufacturers, raising the need to re-execute feature engineering and the construction of new prediction modules. Deep learning can overcome such problems because it does not require feature engineering.

Deep learning is a branch of machine learning, and is a type of algorithm based on feature learning of data. Feature learning is a technique by which features are learned, transforming raw data into a form that can be effectively learned by machine. It avoids issues related to manual feature extraction, allowing the computer to simultaneously learn to use and extract features. This replaces the need for domain-specific subject matter expertise in identifying the important eigenvalues in the research data. Several deep learning architectures have been applied with good results to computer vision, speech recognition, and natural language processing, including deep neural networks (DNN) and recurrent neural networks (RNN).

2.3 Recurrent Neural Networks

Recurrent neural networks (RNNs) are a type of neural network in which each node has a direct one-way connection, and each node has different inputs or outputs at different time points. Each connection has a weight that can adjust the value. The most basic RNN architecture has three nodes which are divided into input, hidden and output nodes. Giles [6] used RNN to make financial exchange rate forecasts. The results indicate that RNN can identify important information within foreign exchange rate data and accurately predict future movements. In addition, the method can extract important feature values from data during the training process.

Long and short-term memory networks (LSTMs) are a special type of RNN that learns long-term dependencies. They were introduced by Hochreiter and Schmidhuber [7], and subsequently modified by many other groups. LSTM is clearly designed to avoid long-term dependencies. LSTM combined with the appropriate gradient-based learning algorithms, can solve the Vanishing or Exploding Gradient Problems that RNN can experience in long-term memory training, and provides a solution for the time series characteristics problem in hard disk failure prediction. Graves [8] used the TIMIT database to evaluate the bidirectional LSTM (BLSTM) and several other network architecture tests in a framewise phoneme classification. They found that bidirectional networks outperform unidirectional networks, while LSTM outperforms RNN, and multilayer perceptrons (MLPs) are much faster and more accurate.

3 Research Methodology

This study was based on data from Backblaze, an online service offering encrypted backups for individuals, families, organizations and businesses in more than 140 countries. As of the end of 2017, Backblaze used 93,240 hard drives ranging in size from 3 TB to 12 TB.

3.1 Research Architecture

As shown in Fig. 1, this study proposes an architecture consisting of three modules: (1) an extraction, cleanup and loading (ECL) module; (2) a deep learning (DL) module; and (3) a prediction and health management (PHM) module. The ECL module accounts for fault mode (FM) and extract and access (ES), while the ML consists of generating modules, and the PHM includes condition monitoring (CM) and a warning system (WS).

3.2 ECL Module

The ECL module consists of two basic components: Failure Mode (FM) and Extraction and Storage (ES).

Backblaze defines a hard disk failure as follows: (1) HDD completely stops working. (2) HDD data indicates imminent failure. When a hard disk meets either of these criteria, it is removed from the server, marked as malfunctioning, and replaced. A “dead” HDD displays no response to read or write commands. Backblaze uses SMART statistics as indicators of imminent disk failure, and a disk determined to be at risk of failure will be removed.

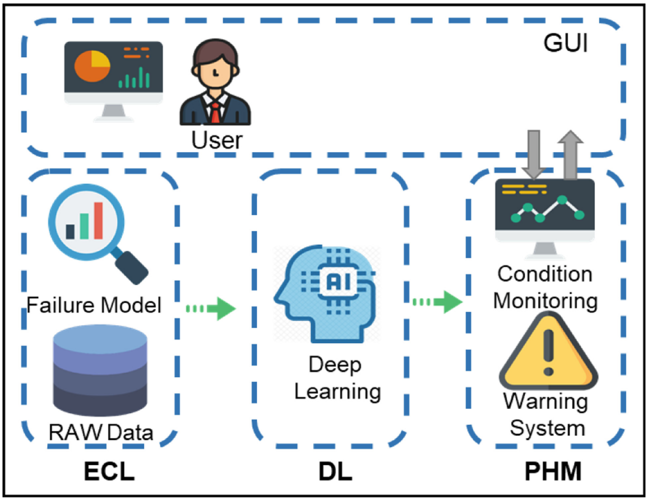

Fig. 1. Research architecture

The extraction and storage components are responsible for collecting SMART data from system log files. Pre-processing includes SMART extraction, noise cleanup and classification. The module performs the following tasks:

- Data extraction – Extracting information from raw data.
- Data cleaning – removal of data noise.
- Data loading – loading data to its final destination.

3.3 DL Module

The second module system consists primarily of a generated model. We define the HDD failure mode and applied SMART training data to implement the deep learning algorithm to solve the predictive monitoring problem for HDD failure. Figure 2 shows an overview of the deep learning model.

In the model construction step, the document system's SMART attributes are already ready for model training. The prediction model for establishing LSTM was built using Keras, a high-level neural network API written in Python that runs as a backend with TensorFlow, CNTK or Theano. Using Keras to build a predictive LSTM network includes three steps: (1) data initialization, (2) fitting it to the training data, and (3) fault prediction. The output is the result of an HDD failure prediction.

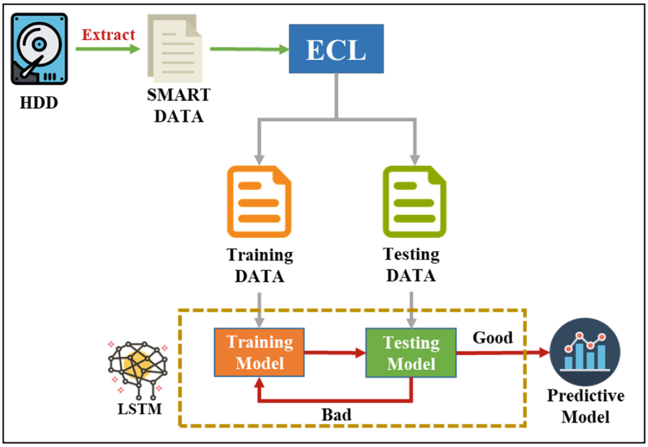

Fig. 2. Deep learning module

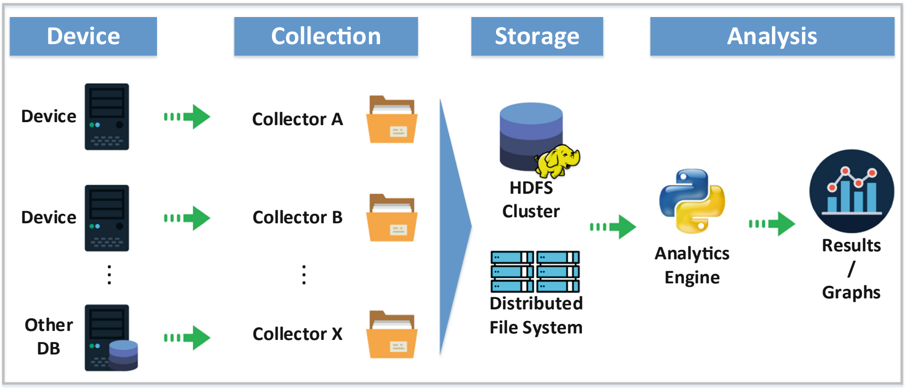

Fig. 3. Research architecture infrastructure concept

3.4 PHM Module

To present the overall research architecture, we propose using a more sophisticated approach for monitoring equipment. Collecting and processing data requires an infrastructure. In Pinheiro's research, the infrastructure describing system health is a large distributed software system that collects and stores hundreds of attribute values from all of Google's servers [9]. It also provides an interface for processing data for any analysis task. We use the health system infrastructure to illustrate the infrastructure of the research architecture, with the infrastructure concept map shown in Fig. 3.

The research architecture infrastructure consists of a data collection layer that collects all historical or streaming data from each device, which is then stored and processed using a distributed file system, and finally processed using an analytics server to product prediction results. Thus far, we have trained our prediction model. To further improve on the original monitoring system, we propose a two-step method as follows: (1) A Value Prediction (VP) component establishes a background service on a device, and the streaming data is sent to the cloud-based analysis platform; and (2) the Warning System (WS) component runs on the analysis server, analyzing the streaming data received from the device and providing device status reports to the user.

To collect SMART data from a hard drive requires a HDD monitoring kit or software. This study used smartmontools (https://www.smartmontools.org/), which consists of two utilities, smartctl and smartd, which monitor SMART data on HDDs and SSDs. This system allows the user to check hard disk SMART data and run various tests to determine disk health. Figure 4 presents a screenshot of smartmontools monitoring SMART data.

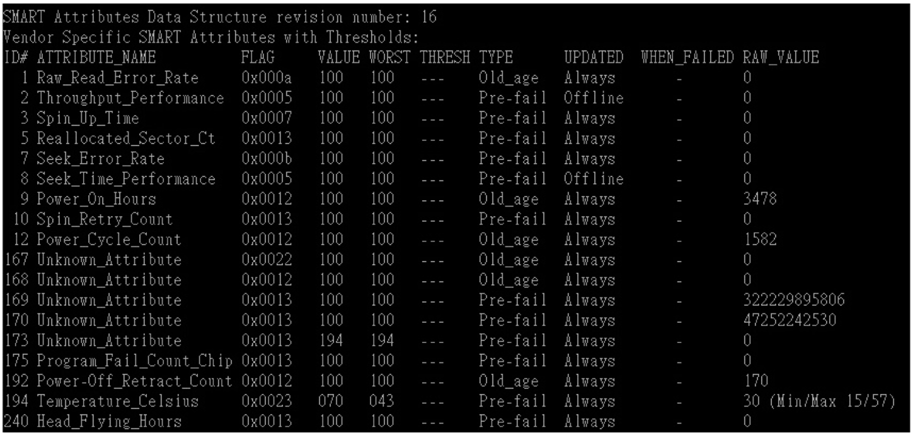

```
SMART Attributes Data Structure revision number: 16
Vendor Specific SMART Attributes with Thresholds:
ID# ATTRIBUTE_NAME          FLAG     VALUE WORST THRESH TYPE      UPDATED  WHEN_FAILED RAW_VALUE
  1 Raw_Read_Error_Rate     0x000a   100   100   ---    Old_age   Always       -       0
  2 Throughput_Performance  0x0005   100   100   ---    Pre-fail  Offline      -       0
  3 Spin_Up_Time            0x0007   100   100   ---    Pre-fail  Always       -       0
  5 Reallocated_Sector_Ct   0x0013   100   100   ---    Pre-fail  Always       -       0
  7 Seek_Error_Rate         0x000b   100   100   ---    Pre-fail  Always       -       0
  8 Seek_Time_Performance   0x0005   100   100   ---    Pre-fail  Offline      -       0
  9 Power_On_Hours          0x0012   100   100   ---    Old_age   Always       -       3478
 10 Spin_Retry_Count        0x0013   100   100   ---    Pre-fail  Always       -       0
 12 Power_Cycle_Count       0x0012   100   100   ---    Old_age   Always       -       1582
167 Unknown_Attribute       0x0022   100   100   ---    Old_age   Always       -       0
168 Unknown_Attribute       0x0012   100   100   ---    Old_age   Always       -       0
169 Unknown_Attribute       0x0013   100   100   ---    Pre-fail  Always       -       322229895806
170 Unknown_Attribute       0x0013   100   100   ---    Pre-fail  Always       -       47252242530
173 Unknown_Attribute       0x0013   194   194   ---    Pre-fail  Always       -       0
175 Program_Fail_Count_Chip 0x0013   100   100   ---    Pre-fail  Always       -       0
192 Power-Off_Retract_Count 0x0012   100   100   ---    Old_age   Always       -       170
194 Temperature_Celsius     0x0023   070   043   ---    Pre-fail  Always       -       30 (Min/Max 15/57)
240 Head_Flying_Hours       0x0013   100   100   ---    Pre-fail  Always       -       0
```

Fig. 4. Smartmontools reading SMART data

The warning system monitors streaming data from the hard disk, with the WS flow chart shown in Fig. 5.

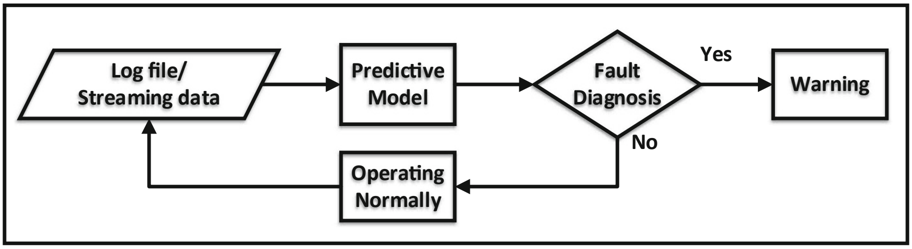

Fig. 5. Warning system flow

Once the hard disk data has been collected from the status monitor, the following steps are taken:

1. Extract the SMART attributes from the streaming data.
2. Perform fault prediction based on modules built by LSTM. If the prediction result is 0, the hard disk is working properly, but a prediction result of 1 indicates imminent damage and triggers a warning message.

4 Conclusion and Future Works

With the emergence of cloud computing and related online services, many large enterprises are increasingly providing services to users through large data centers. Data center maintenance teams must ensure server reliability in order to provide stable, continuous cloud services. In a server, the hard disk is a component particularly prone to failure, thus effective hard disk maintenance can greatly contribute to overall server performance.

The proposed system collects device information for analysis, allowing for the possibility of predicting hard disk failure running the LSTM algorithm on Apache Spark. The system consists of two stages: modeling and real-time prediction. In the modeling stage, historical data is used to generate the LSTM module. Test data is used to verify the ability of the established prediction module to predict hard disk failures. In the real-time prediction stage, data collected from the device in real time is run through the trained prediction module to predict hard disk failure in real-time.

The proposed system allows IT teams to quickly and effectively monitor hard drive condition. Deep learning methods are applied to provide a new kind of insight into data, allowing for improved efficiency in the maintenance of large storage systems. Accurate prediction of device failure reduces the need for scheduled maintenance and ensures that each device is in service for its entire useful life time.

This study proposes an innovative approach to monitoring hard disk failures. The concepts developed extend beyond descriptive statistics and charts. The implementation uses SMART statistics to predict imminent hard disk failure. However, hard disks from various vendors will have different SMART statistics. While each manufacturer follows general guidelines in defining data attributes, the meaning of the SMART values ultimately depends on the definition provided by the hard disk manufacturer. Therefore, this proposed system provides a predictive model for hard disk manufacturers, but the predictive model cannot be applied to hard disks from all other manufacturers. Thus, when predicting and analyzing performance of hard disks from various manufacturers, one must use historical data specific to each disk to establish a new prediction module.

The Internet of Things (IoT) concept uses the Internet as a means of linking physical devices. In the future, IoT will raise many opportunities for improving productivity and efficiency. The cloud analytics platform could be used to collect data from multiple sources and multiple devices, providing significant advantages to managing large numbers of devices in a wide range of domains.

In this research, the BackBlaze data of the hard disk is used to train the module that can predict the hard disk failure based on the SMART data and issue alert when there is a possibility of failure. In order to improve the results of the prediction, we suggest developing a complete solution, which can predict the timing of the hard disk failure more accurately. Furthermore, it can inform the user when the monitored hard disk is about to be failed and show the remaining lifetime, helping the user understand the status of the hard disk more clearly.

References

1. Andy, K.: Backblaze Hard Drive Stats for 2017 (2017). https://www.backblaze.com/blog/hard-drive-stats-for-2017/
2. Canizo, M., Onieva, E., Conde, A., Charramendieta, S., Trujillo, S.: Real-time predictive maintenance for wind turbines using big data frameworks. In: 2017 IEEE International Conference on Prognostics and Health Management (ICPHM), pp. 70–77. IEEE, June 2017
3. Zhao, P., Kurihara, M., Tanaka, J., Noda, T., Chikuma, S., Suzuki, T.: Advanced correlation-based anomaly detection method for predictive maintenance. In: 2017 IEEE International Conference on Prognostics and Health Management (ICPHM), pp. 78–83. IEEE, June 2017
4. Samuel, A.L.: Some studies in machine learning using the game of checkers. IBM J. Res. Dev. **3**(3), 210–229 (1959)
5. Brownlee, J.: Discover Feature Engineering, How to Engineer Features and How to Get Good at It (2014). https://machinelearningmastery.com/discover-feature-engineering-how-to-engineer-features-and-how-to-get-good-at-it/
6. Giles, C.L., Lawrence, S., Tsoi, A.-C.: Rule inference for financial prediction using recurrent neural networks. In: IEEE Conference on Computational Intelligence for Financial Engineering, p. 253. IEEE Press (1997)
7. Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural Comput. **9**(8), 1735–1780 (1997)
8. Graves, A., Schmidhuber, J.: Framewise phoneme classification with bidirectional LSTM and other neural network architectures. Neural Networks **18**(5–6), 602–610 (2005)
9. Pinheiro, E., Weber, W.-D., Barroso, L.A.: Failure Trends in a Large Disk Drive Population (2007)