IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received June 21, 2018, accepted July 29, 2018, date of publication August 21, 2018, date of current version September 21, 2018.

Digital Object Identifier 10.1109/ACCESS.2018.2866522

A Large-Scale Study of I/O Workload's Impact on Disk Failure

SONG WU[1,2], YUSHENG YI[1,2], JIANG XIAO[1,2], HAI JIN[1,2], AND MAO YE[3]

[1]Services Computing Technology and System Lab, School of Computer Science and Technology, Huazhong University of Science and Technology, Wuhan 430074, China
[2]Cluster and Grid Computing Lab, School of Computer Science and Technology, Huazhong University of Science and Technology, Wuhan 430074, China
[3]Tencent Inc., Shenzhen 518000, China

Corresponding author: Jiang Xiao (jiangxiao@hust.edu.cn)

This work was supported in part by the National Key Research and Development Program under Grant 2016YFB1000501 and in part by the National Science Foundation of China under Grants 61472151 and 61702203.

ABSTRACT In large-scale data centers, disk failure is the norm rather than an exception. Frequent disk failure noticeably hurts user experience and results in unavailability of data in the worst case. Previous researches from both industry and academia have studied the reasons of disk failure; however, there is a lack of knowledge of the intrinsic relation between failed disks and their I/O workload. In this paper, we collect and investigate about four billion drive hours I/O traces over 500 000 disks in Tencent's data centers. Our focus is to first exploit the key characteristics of I/O workload that influences disk reliability. We further present the impact of these I/O workload features on lifespan of disks and uncover the root causes. Finally, we introduce a new metric to accurately identify the "dangerous" I/O workload which is extremely harmful to disk health. To the best of our knowledge, this research is by far the first in-depth analysis of the I/O workload's impact on disk reliability and opens up a new dimension for I/O schedule policy in data centers.

INDEX TERMS Disk failure, I/O workload, duty cycle, bandwidth.

I. INTRODUCTION

Disk drives have become the dominated storage medium in large-scale data centers as it can provide substantial capacity with low price. But it also incurs huge maintenance cost due to high failure rate and massive amounts of population. In production data centers, the probability of disk replacement has been much higher than replacing other components [1], [2]. Furthermore, since the number of disk drives is usually more than the number of other components (e.g., CPU and memory) in enterprise servers, the number of vulnerable drives is much more than other components as well. Consequently, disk failure has become a crucial issue in data centers and it is of great importance to figure out its root cause.

Prior studies on disk failure can be classified into two groups. The first group concentrates on the inherent characteristics of disk drives that reflect the complicated design and implementation, such as disk head [3], media material [3], and interface [4], etc. The second group is the environment factors, such as temperature [2], [5], humidity [5], air quality [5], and I/O workload [4], [6], etc. Unlike the inherent characteristics belonging to default setting by disk vendors, the environment factors are more convenient to be adjusted for storage reliability enhancement and thus have received increasing interest from data centers.

Among the environment factors, I/O workload is unique for its complexity and granularity. First, I/O workload has richer properties than the other environment factors, e.g., the amount of transferred data, I/O request types, and I/O request patterns. Such complexity serves as a huge challenge to understand the relation with disk failure. A pioneer work from Google [6] concludes that I/O workload, represented by the average bandwidth per disk, is weakly related to disk failure. However, it is unlikely to show us the whole picture since it only takes one aspect of I/O workload into account, which motivates us to conduct a more comprehensive study. Second, the collection of real-world I/O workload in production data centers entails considerable impediment for further investigation. Specifically, it is required to collect I/O trace at disk granularity, which consumes more time than gathering the coarse-grained (e.g., rack-level or data center level) temperature or air quality. Even worse, I/O workload changes at different stages of an application, which requires a more frequent collection and a larger data repository than other environment factors. Thus, current data centers are reluctant to monitor and gather such historical I/O workload trace. Due to aforementioned two reasons, there have been very limited studies to present in-depth analysis of how I/O workload influences disk failure.

To that extent, we first collect more than 4 billion drive hours I/O traces over 500,000 disks in Tencent's production data centers. For each disk, we collect duty cycle and bandwidth in each five minutes interval.

We then exploit and demonstrate the subtle features of I/O workload from both temporal and spatial dimensions, respectively. For the former one, we propose the *average duty cycle* to describe the busy level of a disk and the *ratio of large duty cycle* to denote the workload burst within the five minutes interval. To further explore how much data a disk transfers and the average ratio of writing data to these transferred data, the corresponding spatial features including the *average bandwidth* and the *average ratio of write bandwidth* are presented. Moreover, we introduce a new metric, called *average intensity of sequential request* (*AISR*) to accurately and comprehensively capture the impact of I/O workload on disk failure. Our key observations are listed as below:

- There is a clear trend showing that larger *average duty cycle* is associated with higher *annual failure rates* (*AFR*). A large *average duty cycle* (greater than 50%) can increase *AFR* by as much as 347%. To achieve a low *AFR*, a simple strategy is to limit all duty cycle less than 50%. To make it practical, we study the relationship between *average duty cycle* and *AFR* and then propose a flexible strategy to allow infrequent occurrence of large duty cycle. Besides, the relationship between *average duty cycle* and *AFR* implies that roughly increasing disk utilization is probably not a suitable way to reduce the *total cost of ownership* (TCO) of data centers, which is quite different from CPU resource management.
- The impact of *average bandwidth* on *AFR* can be divided into two parts: first, when *average bandwidth* is small (e.g., less than 5000 KB/s), no obvious correlation with *AFR* is found; Second, while *average bandwidth* is large (e.g., greater than 5000 KB/s), it shows a negative correlation with *AFR*. The reason lying behind is that *average bandwidth* not only represents the speed of data transfer but also indirectly indicates how many random requests a disk serves. Disks transfer data slowly in proportion as they serve random requests. Since the random request introduces more mechanical movement for track seeking than sequential request, it accelerates the disk wear-out.
- To quantify the randomness of I/O workload, we propose a new metric *AISR*. Our experiments show that there is a significant correlation between *AISR* and *AFR*. With *AISR*, we can not only measure *AFR* for disk drives, but also provide a new perspective for I/O scheduling in data centers.

II. TRACE COLLECTION AND BACKGROUND

A. DATA SOURCE

The Chinese Internet giant Tencent builds data centers around the world that serve billions of people (963,000,000 monthly active Wechat users in 2017Q2 [7]). We collect traces from 58 of these data centers for one year. The trace contains 9,177 disk failures and more than 4 billion disk hours data of I/O workload from 530,522 disk drives.

Tencent's data centers are built and maintained in a typical way, that covers the majority of server configuration in production data centers. More specifically, enterprise servers are used to provide services like instant message, video games, video on demand, and cloud, etc. Each server is equipped with a dual-core or quad-core CPU processor, 2G-128G memory, and multiple SATA2/SATA3 disk drives for various application demands. Hence, we are reasonable to say that our conclusions and strategies of workload assignment are representative and generally suitable for most of modern data centers.

B. DISK I/O WORKLOAD

Tencent deploys a probe daemon on each server to monitor the major components, including CPU, memory, network, and storage. The probe daemon collects component status by monitoring the resource usage of components from the */proc* pseudo-filesystem [8]. It contains a hierarchy of special files that represent the current state of different server components. Based on the component status, the probe detects anomalies and raises warnings timely for each server. In this way, Tencent builds up a large run-time monitoring system to manage and maintain data center servers.

By monitoring the file */proc/diskstats*, the probe daemon collects 11 attributes for storage. To avoid high overhead of storage space and network bandwidth, Tencent allows us to store three of them in a five minutes interval, including the bandwidth of reading (KB/s), the bandwidth of writing (KB/s), and the utilization of disk (%). The commonly used workload metric IOPS is evaluated by the combination of the duty cycle and the bandwidth in Section IV. To our best knowledge, it is the first trial to correlate workload and disk failure in such a fine-grained interval. Since duty cycle and bandwidth can measure whether a disk experiences a heavy workload, we refer them as workload intensity.

Based on the duty cycle and the bandwidth, we proposed six features to study the relationship between I/O workload and disk failure. We list them in Table 1 as well as some frequently used notations in this paper.

C. DISK FAILURE

To prevent server from storage system failure, data centers simply replace any suspected components and pass them to the vendor for further diagnosis, especially when the components are under warranty. This conservative strategy is prone to achieve an exaggerated number of failed drives and an overestimated *AFR*. Focusing on diagnosing and repairing

TABLE 1. Notation table.

Symbol	Description
d	duty cycle, i.e., ratio that a drive is active out of unit time
b	bandwidth, i.e., number of bytes reading/writing per second
b^w	number of bytes writing per second
t	threshold to recognize a large duty cycle
D_{avg}	average duty cycle
D_l	ratio of large duty cycle
D_{cv}	coefficient variation of duty cycle
B_{avg}	average bandwidth
B_w	percentage of write bandwidth
AISR	average intensity of sequential request
AFR	annual failure rate

returned drives, disk vendors report that 20%-43% of the returned disks are healthy [1], [9]. Therefore, an all-around examination of disk failure should be done to present an accurate number of failed disks. To this end, Tencent develops a three-step disk failure detection with the assistance of disk vendors. Disk drives are divided into four groups based on their healthy status (i.e., healthy, abnormal, error, failure).

At the beginning, all disk drives stay in a healthy phase. The probe daemon, introduced in the last subsection, parses the file */var/log/dmesg* and examines whether there exist lines containing keywords like 'I/O error' or 'disk read-only'. Such logs lead to a conversion from a healthy disk to an abnormal disk, which equals to the conservative manner as discussed above. To avoid treating the storage failure caused by other components of storage subsystem as a disk failure, a light-weighted self-inspection is launched. It examines whether the common operations like read and write function well, and whether the performance of the drive keeps in a good condition. We recognize an abnormal drive failed any test in the self-inspection as an error drive, which is delivered to an error component repository. For most of IT companies, self-inspection is usually the last step to confirm a failed disk.

In Tencent, a rigorous offline checking and recovery process rechecks the error drive in a customized environment, developed by Tencent cooperated with disk vendors. The *Self-Monitoring, Analysis and Reporting Technology* (SMART) and other vendor-specific attributes are inspected from the perspective of vendors. An error drive failed to pass the offline check is recognized as a failed disk. From the vendor's report, the percentage of healthy returned disks can be reduced from 40% to 10% by enabling the offline checking and recovery process. In this study, we use the failed disk to calculate *AFR*.

D. OTHER FACTORS

1) TEMPERATURE, HUMIDITY AND AIR QUALITY

Several studies have investigated the influence of temperature and humidity on disk failure in data centers [2], [5], [10], [11]. They divide data centers into free cooling and non-free-cooled. The former, exposing components to aggressive environment conditions (e.g., low temperature and higher relative humidity), are supposed to address the tradeoffs between the cooling energy and the hardware component reliability. Data centers in our study are consistent with the latter, where air conditions control the temperature, the humidity, and the air quality globally to offer disk drives a proper and stable environment. Hence, these environment factors do not become major determinants of disk failure in our study.

2) DISK MODELS, MANUFACTURERS AND VINTAGES

New disks are constantly introduced in Tencent data centers every year, which constructs a multi-model population mix of disk drives. We focus on ten disk models, occupying 93.79% of the population. Although some disk models show a slightly higher *AFR* than others, the *AFR* trends are similar across models. This confirms Googles' conclusion [6] that the mix population of disk products does not change impacts of environment factors on disk failure. Therefore, we do not show a breakdown of drives per manufacturer, model, or vintage, but treat them as a whole to represent a typical situation of modern data centers.

3) AGE AND SMART

Age is one of the most important factors to analyze the reliability of disk drives [1], [6]. SMART is a commonly used data to evaluate disk health status, which provides a good disk failure prediction performance in [12] and [13]. However, we argue that they are not the direct factors affecting disk failure but indicators of disk aging. Because a serious aging of disk leads to a disk failure, metrics of disk aging are good on failure analysis and failure prediction. Different from disk aging measurement, our target is to figure out the direct factors of disk aging. Thus, we start from the inherent factors and environment factors, especially the assigned jobs (e.g., workload) and the working environment (e.g., temperature and humidity). As mentioned above, Tencent provides a stable environment for servers in their data centers. Thus, we focus on the workload to expose the reason and pattern of disk aging. Following this, we can combine features of I/O workload with age or SMART to make failure analysis and to predict disk failures which may be more practical for improving disk reliability actual situations in data centers.

III. TEMPORAL AND SPATIAL CHARACTERISTICS OF WORKLOAD

In this section, we investigate the characteristics of I/O workload to explore their impacts on disk reliability. We first quantify how busy a disk is with two temporal features of I/O workload (D_{avg} and D_l) in different granularities. Then, we present two spacial features of I/O workload (B_{avg} and B_w) to explore how much data a disk transfers and the average ratio of writing data to these transferred data.

A. TEMPORAL FEATURES OF WORKLOAD

As a coarse-grained metric, D_{avg} quantifies the average working hours during a disk's life and shows a high relevancy to disk failure. D_l is a fine-grained representation of disk working hours, which answers the question that how many large duty cycles can affect the disk failure significantly. With D_{avg} and D_l, we propose two strategies to avoid disks from a high *AFR*.

1) AVERAGE DUTY CYCLE

For each disk, we average all duty cycles as D_{avg}, as shown in Formula 1. D_{avg} less than 1% means a disk is mounted on a server with little storage demand. D_{avg} of 100% indicates a disk is always busy in its lifetime. To give an overview of the busy level of disks in data centers, we present the distribution of D_{avg} with an increment of 5% in Figure 1. Figure 2 shows the impact of D_{avg} on *AFR* with deepened bar to distinguish high *AFR* (greater than 5%) and low *AFR* (less than 5%). Besides, we use a red dashed line to fit the trend of *AFR* by the method of local polynomial regression fitting. We add the deepened bar and the red dashed line for each *AFR* figure in the following sections.

$$D_{avg} = \frac{1}{n}\sum_{i=1}^{n} d_i \tag{1}$$

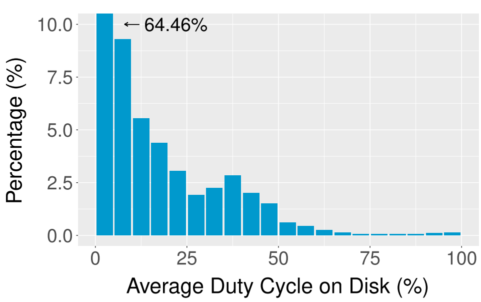

FIGURE 1. **Distribution of *Average duty cycle*. More than 98% disk drives have an *Average duty cycle* less than 50%. We truncate the tallest bar to show the detail of all bins.**

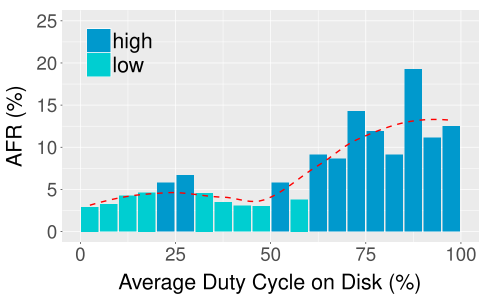

FIGURE 2. ***AFR* vs. *Average duty cycle*. It shows *Average duty cycle* has an obvious impact on *AFR*, especially when *Average duty cycle* is greater than 50%. We deepen the bar to emphasize a high *AFR* over 5% and plot a red dash line to fit the trend of *AFR*.**

Figure 1 shows that D_{avg} is relatively low among most disk drives and only less than 2% of disk drives have an D_{avg} greater than 50%. It indicates that disk drives with a low D_{avg} is the norm in data centers. We explain this with two reasons. The first one is the diversity of applications in data centers. Not all servers are dedicated to storage. Many of them are employed for jobs like connection and computation. Naturally they serve few I/O workload and thus exhibit a small D_{avg}. The second reason is to meet the *service level agreements* (SLAs). Data centers are provisioned for peak load, which is significantly higher than average load. This leads to a considerable resource idleness on average and generates a small D_{avg}. The common situations are also reported by Kaplan *et al.* [14] that the average data center resource utilization is low.

Consistent with our expectation, Figure 2 shows that D_{avg} has a significant impact on *AFR*. More specifically, *AFR* is low and increases slowly when D_{avg} is less than 50%. While D_{avg} is greater than 50%, *AFR* bursts from 3.05% ($D_{avg} = 45\%$) to 19.23% ($D_{avg} = 85\%$). In general, *AFR* of disks with D_{avg} less than 50% is 3.47 times as much as *AFR* of disks with D_{avg} greater than 50%, which indicates keeping D_{avg} under 50% is a good way to improve disk reliability. Although disks with ADC of 30% has a higher AFR than other disks with ADC less than 50%, they are not the main contributor of AFR improvement. We focus on disks with ADC greater than 50%. A simple strategy is to restrict all duty cycles less than 50%. However, many large data centers are trying to keep a high-level utilization of servers to improve TCO. From our result, improving TCO by roughly increasing utilization is probably not a suitable way for disk drives. Though higher disk utilization (i.e., D_{avg}) brings more profit for data centers, it also results in more disk failures that decrease the average lifespan of disk drives. Thus, there is a trade-off between the benefit of high disk utilization and the expensive recovery cost of failed disks to be made for data centers.

Finding 1: Disk drives spending more than half of their lifetime on I/O workload exhibit much higher (3.47 times) AFR than the rest. Different from CPU resource management, roughly increasing disk utilization is probably not a suitable way to reduce TCO of data centers. Thus, the trade-off between the benefit of high disk utilization and the expensive recovery cost of failed disks still leaves as an open issue.

2) RATIO OF LARGE DUTY CYCLE

For a better understanding of duty cycle, especially the large duty cycle, we present a fine-grained feature D_l to answer the question that how many large duty cycles will affect *AFR* significantly. In this part, we present loose strategy to allow occurrence of large duty cycles and give a suggestion to decrease *AFR* for disks with heavy workload.

For each disk, D_l is evaluated by the ratio of number of large duty cycles to the total number of duty cycles. We first choose a threshold (t) of duty cycle (from 0% to 100%) to distinguish whether a duty cycle is a large duty cycle in Formula 2. The t determines the measurement standard of large duty cycle for all disk drives. Then, we calculate D_l for each disk under the special standard in Formula 3. Note that D_l depends on the t, a strict (large) t reduces D_l of disk drives.

$$f(d_i) = \begin{cases} 1, & d_i \geq t \\ 0, & d_i < t \end{cases} \tag{2}$$

$$D_l = \frac{\sum_{i=1}^{n} f(d_i)}{n} \tag{3}$$

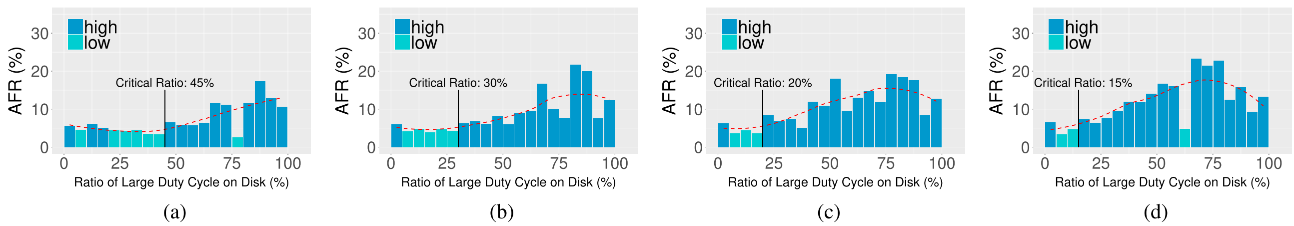

FIGURE 3. ***AFR*** **vs.** ***Ratio of large duty cycle*** **under four thresholds. A duty cycle is recognized as a large duty cycle if it is greater than the threshold. In each graphs,** ***AFR*** **trend can be divided into three phases: a phase of low** ***AFR*****, a phase of** ***AFR*** **increasing, and a phase of** ***AFR*** **reduction. Threshold affects the length of them. (a) Large duty cycle threshold = 50%. (b) Large duty cycle threshold = 60%. (c) Large duty cycle threshold = 70%. (d) Large duty cycle threshold = 80%.**

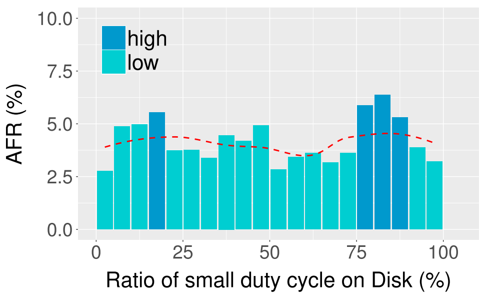

FIGURE 4. ***AFR*** **vs. Ratio of small duty cycle (small duty cycle threshold = 5%). The ratio of small duty cycle is weakly related to disk failure.**

We set four ts at duty cycle of 50%, 60%, 70%, and 80% and plot figures under these ts to study the impact of D_l on *AFR* in Figure 3. It shows that *AFR* grows with D_l consistently in all figures. More specifically, every *AFR* trend can be divided into three phases. They are the low *AFR* phase, the increasing *AFR* phase, and the *AFR* reduction phase. t affects the length of each phases. Across the four figures, we find smaller t corresponds to a longer low *AFR* phase. For example, the low *AFR* phase lasts when D_l is less than 45% in Figure 3(a). While t grows from 50% to 60%, the low *AFR* phase only lasts when D_l is less than 30% in Figure 3(b). For a low *AFR*, disk drives should stay in the first phase. We thus define the end of the low *AFR* phase as the critical D_l, shown in each figure. It is the key of our new strategy that breaks the limit of duty cycle less than 50% mentioned in last part.

The fundamental principle of our new strategy is to make sure D_l is less than critical D_l for each disk under each t. For example, we allow the appearance of duty cycle greater than 80% ($t = 80\%$, corresponding to Figure 3(d)), but its corresponding D_l should be less than 15% (critical $D_l = 15\%$). To make this strategy practical, data center should check whether critical D_l is unsurpassed for each disk and limit the occurrence of large duty cycles leading to an impending break of critical D_l. In this way, we not only allow the occurrence of duty cycle greater than 50%, but also maintain disk drives at a low *AFR*.

Except for the critical D_l, there are two noteworthy problems to study. The first one is a mirror question that whether the number of small duty cycles has a positive impact on *AFR* reduction. Similar to large duty cycle, we use a threshold t' to recognize small duty cycles and a ratio to represent the number of duty cycles less than t' to the number of duty cycles for each disk. We plot the relationship between the ratio and *AFR* in Figure 4 with a t' of 5%. We also observe results with t's at 10%, 20%, and 30%. But no figure shows a significant variation or obvious trend on *AFR*. It indicates that the number of small duty cycles is weakly related to *AFR*. Thus, data centers only need to focus on the occurrence of large duty cycles.

The second problem is the reduction phase in Figure 3. We note that *AFR* begins to decrease when D_l is greater than 85%, 80%, and 75% in Figure 3(b), 3(c), and 3(d) respectively. Since larger D_l means more time a disk spends on a busy status, the *AFR* reduction indicates disk drives 'always busy' are healthier than disk drives 'almost always busy', especially for disk with heavy workload. Based on the definition of D_l, the 'almost always busy' drives have more duty cycles less than the t than the 'always busy' drives, indicating a more unstable and looser distribution of duty cycles. Thus, we speculate *AFR* is sensitive to the stability of workload so that the workload switching between large duty cycle and small duty cycle accelerates the wear-out of disk drives.

$$D_{cv} = \frac{\sqrt{\sum_{i=1}^{n}(d_i - D_{avg})^2}}{D_{avg}} \tag{4}$$

To verify our speculation, we present the stability of duty cycle by the *coefficient of variation* (D_{cv}) in Formula 4. D_{cv} is the ratio of standard deviation of duty cycles to the mean of duty cycles (D_{avg}) for each disk drive, as suits for comparing vibration of samples with different averages like D_{avg}. We plot the impact of D_{cv} on *AFR* in Figure 5. In this figure, *AFR* is low and stable when D_{cv} is less than 0.5. While D_{cv} is greater than 0.5, *AFR* grows rapidly. Table 2 lists D_{cv} of disk drives under t and D_l. We highlight the inflection point where D_{cv} begins to get down. Table 2 shows the corresponding relationship between the *AFR* reduction in Figure 5 and the decreasing of D_{cv} clearly. For example, *AFR* decreases when D_l grows from 75% to 100% in Figure 3(d) ($t = 80\%$) corresponds to D_{cv} reduces from 0.721 to 0.453 in Table 2. Therefore, we are reasonably to attribute the *AFR* reduction phase in Figure 3 to the growing stability. To improve disk

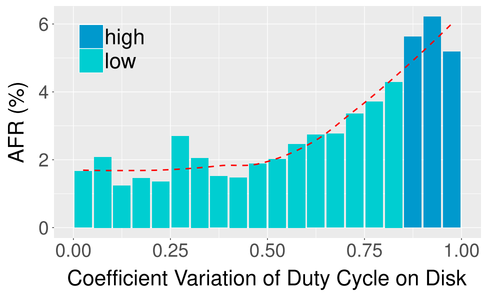

FIGURE 5. ***AFR* vs. *Coefficient variation of duty cycle*. *AFR* grows rapidly when *Coefficient variation of duty cycle* is greater than 0.5.**

TABLE 2. **Coefficient variation of duty cycle.**

threshold (t) \ ratio (D_l)	60%	70%	80%	90%
60%	0.568	0.672	0.732	0.502
70%	0.632	0.707	0.592	0.434
80%	0.734	0.721	0.586	0.453

reliability, we suggest a stable I/O workload for disk drives, especially for those with a heavy workload.

Finding 2: Large duty cycle exhibits a subtle influence on AFR. On one hand, it shows a positive relation with AFR in general, and restricting the frequency of its appearance can effectively prevent disk failure. On the other hand, there exists a counter-intuitive phenomenon that the disks with saturated large duty cycle exhibit lower AFR than those with unsaturated large duty cycle, which implies that less workload vibration can make disks healthier, especially for those with a heavy workload.

B. SPACIAL FEATURES OF WORKLOAD

In this part, we present two spacial features B_{avg} and B_w. B_{avg} is used to explore the impact of the average amount of data transferred between disk and memory during a disk's lifetime on disk failure. B_w presents the average ratio of writing data to the transferred data to show the impact of I/O request type on disk failure.

1) AVERAGE BANDWIDTH

We first study the average amount of data transfer by B_{avg}, shown in Formula 5. Due to disk fragmentation and low spin rate, B_{avg} is exceptionally hard to achieve the ideal upper bound of bandwidth in production environments and disk drives thus are highly concentrated at a low B_{avg}. In our study, only 0.34% disk drives have an B_{avg} greater than 9000 KB/s. To avoid a long-tail distribution, we divide disk drives of B_{avg} from 0 KB/s to 9000 KB/s into 20 bins and add disks with B_{avg} greater than 9000 KB/s in the last bin.

$$B_{avg} = \frac{1}{n}\sum_{i=1}^{n} b_i \tag{5}$$

We first present the distribution of B_{avg} in Figure 6. Similar to D_{avg}, most of the disk drives accumulate at a low B_{avg}. The first bar contains more than half disk drives (59.12%). More specifically, a majority of them (i.e., 24.3% of all disk drives) has B_{avg} less than 20 KB/s. The possible reason is that the workload on these very low B_{avg} disks is not an application workload but an intermittent workload of system logs or component status recording like our metrics, as generates a small B_{avg}. The distribution of B_{avg} indicates there exists enough space to support the growing workload intensity.

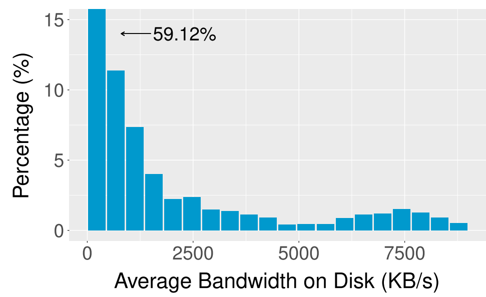

FIGURE 6. **Distribution of *Average bandwidth*. Only a few of disk drives transfer data at a high speed. To limit the long tail of *Average bandwidth*, we truncate the *Average bandwidth* at 9000KB/s.**

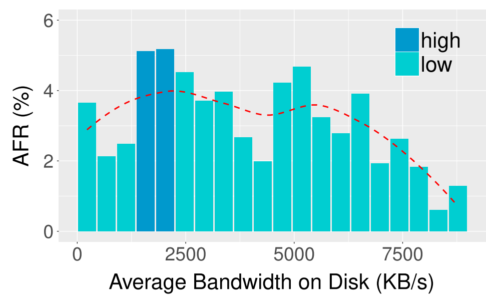

FIGURE 7. ***AFR* vs. *Average bandwidth*. No obvious *AFR* trend is found when *Average bandwidth* is less than 5000 KB/s. While *Average bandwidth* is greater than 5000 KB/s, it shows an unexpectedly negative correlation between them.**

Then, we present the impact of B_{avg} on AFR in Figure 7. Different from D_{avg}, we do not observe a growing AFR. When B_{avg} is less than 5000 KB/s, AFR vibrates but the red dash line, representing the AFR trend, stays at 3.5%. It implies that B_{avg} is not a critical factor of disk reliability, which is consistent with [6]. However, when B_{avg} is greater than 5000 KB/s, it shows an unexpectedly negative correlation between AFR and B_{avg}, which is also emphasized by the Pearson's correlation coefficient of −0.537. It indicates disk drive of dedicated storage server is healthier than those transferring less data. We explain the counter-intuitive phenomenon in Section IV.

Finding 3: B_{avg} is not a predominant reason for disk failure, but it can also be used as an indicator of disk failure under special situation. That is, in the case of heavy I/O workload, the disks with lower B_{avg} suffer higher risks than the ones with higher B_{avg}. This phenomenon indicates that the main factor affecting disk failure is the pattern, instead of the amount of data transfer to disk drives.

2) AVERAGE RATIO OF WRITE BANDWIDTH

We then study I/O request types (read and write) by B_w in Formula 6. We define B_w as the average ratio of write bandwidth to the total bandwidth (including write bandwidth and read bandwidth) for each disk. A large B_w means a drive transfers more data in disk than out of disk.

$$B_w = \frac{1}{n}\sum_{i=1}^{n}\frac{b_i^w}{b_i} \tag{6}$$

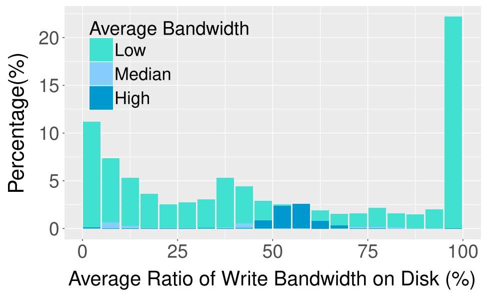

FIGURE 8. **Distribution of *Average ratio of write bandwidth* under three *Average bandwidth* levels (0-3000 KB/s, 3000-6000 KB/s, and 6000-9000 KB/s) to give an overview of speed and direction of data transfer in production data centers. Most disk drives concentrate on the ends of figure. For disks with a high *Average bandwidth*, their *Average ratio of write bandwidth* are close to 55%, representing a balanced bandwidth of reading and writing.**

We first depict the distribution of B_w in Figure 8 under three B_{avg} levels to provide an overview of speed and direction of data transfer in production data centers. Firstly, we observe disks accumulate at the low end and the high end, especially the high end containing 22.23% of disk drives. We contribute it to the fact that most disks are not mounted on a dedicated storage server. Thus, they are easily skewing to the ends since single type of I/O request can easily dominate the I/O request queue, which is idle for most of time. Secondly, disk drives with large B_{avg} are all distributed close to B_w of 55%, which indicates balanced bandwidth is a typical bandwidth for server with heavy workload. Thus, the two typical servers must be taken into account when we need to simulate disk I/O workload in production data centers. One typical server is used for testing lightweight workload, dominated by a single I/O workload type. The other tests a heavy workload consisting of balanced reading workload and writing workload.

Then, we study the impact of B_w on *AFR* in Figure 9. Different from B_{avg}, we find B_w has an obvious influence on *AFR*. When B_w is less than 30%, *AFR* grows linearly with B_{avg} with a correlation coefficient of 0.78. With a B_w greater than 30%, *AFR* grows slowly. Figure 9 indicates reading is healthier to disk drives than writing. We attribute the reason to the in-disk cache, which is a local cache to accelerate disk. When OS sends a reading request to disk, the disk first checks whether the requested data exists in the in-disk cache. If cache hits, the only action is to send data from the cache to memory. In this way, the reading request completes without any mechanical movements within disk

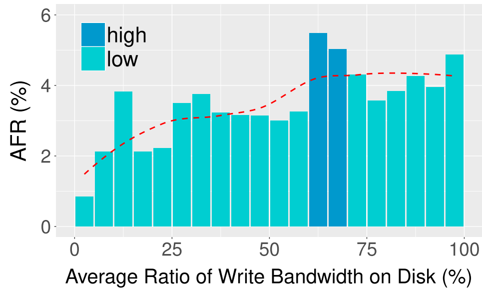

FIGURE 9. ***AFR* vs. *Average ratio of write bandwidth*. The growing *AFR* indicates disk reading is healthier to disk reliability than disk writing.**

drives. Though the on-disk cache can also be a staging buffer for some special workload with high writing locality, writing requests inevitably results in a mechanical movements in disk drives. Thus, writing request brings more wear-out on disks than reading request.

Finding 4: Write operations bring more fatigue to disk drives than read operations. The disk drives with more reads have smaller AFR. It is mainly derived from in-disk cache mechanism which reduces actual mechanical movement of disks.

C. SUMMARY

In this section, we explore temporal features and spacial features of workload to study their impact on disk reliability. The current results of our experiments show that D_{avg} and B_{avg}, describing lifetime workload intensity for disk drives, affect *AFR* in different ways. Despite these interesting findings, it is plausible to explain many noticeable details. For example, the amount of transferred data, indicating the amount of jobs done by drives, should be related to wear-out of disk drive. But in our study, the large number of data transfer makes disk drives healthier than those with less data transfer. Therefore, it is crucial to uncover the relationship between these two features for a comprehensive study. To the best of our knowledge, no such work exists in the literature.

IV. RANDOMNESS OF I/O WORKLOAD

This section discusses the complex relationship of D_{avg} and B_{avg} that describes an intrinsic nature of disk I/O workload, known as randomness. To explicitly represent it, we make a quantitative analysis by introducing a novel metric *AISR*. This can help determine the influence of I/O workload on disk reliability in a fine-grained manner.

A. INTRODUCTION OF WORKLOAD RANDOMNESS

We first present the relationship between D_{avg} and B_{avg} in Figure 10. In this figure, we use the same bins as the D_{avg} distribution in Figure 1 and divide disk drives in each bin by B_{avg}. Since the maximum B_{avg} has a significant variability across different D_{avg} bins, disk drives are divided by the quantile of B_{avg} (10%, 30%, 50%, 70%, and 90%). We use a red line to represent the average B_{avg} in each D_{avg} bin. In this

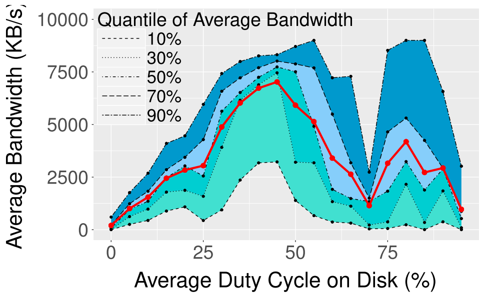

FIGURE 10. Distribution of ***Average bandwidth*** within ***Average duty cycle*** bins. We first plot points of ***Average bandwidth*** quantile (10%, 30%, 50%, 70% and 90%) vertically to show ***Average bandwidth*** distribution in each ***Average duty cycle*** bin. Then, we connect the points horizontally to display the distribution changing as ***Average duty cycle*** grows. The red line represents the average ***Average bandwidth*** in each ***Average duty cycle*** bin. The ***Average bandwidth*** is positively correlated to ***Average duty cycle*** when ***Average duty cycle*** is less than 50%. While ***Average duty cycle*** grows over 50%, the falling ***Average bandwidth*** indicates a weakly negative correlation between ***Average duty cycle*** and ***Average bandwidth***.

figure, we find the linear relationship between D_{avg} and B_{avg} only exists when D_{avg} is less than 50%. When D_{avg} is greater than 50%, B_{avg} decreases with it. We call these disk drives aggregating at the area of high D_{avg} as an 'abnormal' disk. They are likely to have a poor performance since a large D_{avg} and a small B_{avg} means they spend more time on transfer the same amount of data than other disks.

Then, we present a combined *AFR* to find the way D_{avg} and B_{avg} affect *AFR* simultaneously in Figure 11. We use the same bins as the impact of D_{avg} on *AFR* in Figure 2 and divide disk drives in each bin by B_{avg} to find the difference between the combined impact to the individual impact from D_{avg} and *AFR*, respectively. To avoid significant variability across different D_{avg} bins, we use quantile of B_{avg} (33% and 66%) to divide disk drives into three B_{avg} levels in each ABC bin. Different lines correspond to impact from different B_{avg} levels.

In Figure 11, *AFR* of all three lines are similar. They are consistent with the original *AFR* trend of Figure 2, indicating a strong impact of D_{avg} on *AFR* across all B_{avg} levels.

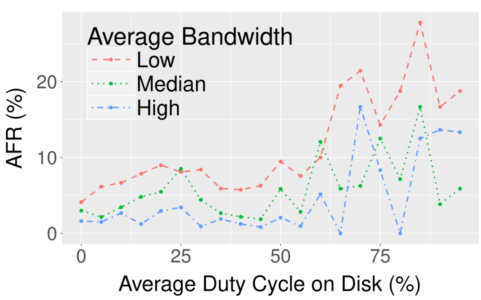

FIGURE 11. ***AFR*** vs. ***Average duty cycle*** under three ***Average bandwidth*** levels. To avoid significant variability across different ***Average duty cycle*** bins, we use the quantile of ***Average bandwidth*** (33% and 66%) to divide disks. We calculate and plot ***AFR*** for each ***Average bandwidth*** level and each ***Average duty cycle*** bin to show the combined impact of them.

For B_{avg}, the almost disjoint three lines of *AFR* trend is consistent with the *AFR* reduction of Figure 7, indicating B_{avg} does affect disk reliability in a way. From this figure, another interesting observation is that the most vulnerable drives, which have a large D_{avg} and a small B_{avg}, match the characteristics of 'abnormal drives'. This coincidence leads us to explore the 'abnormal drives'.

To understand the 'abnormal disk', we begin with its workload characteristics. Large D_{avg} and small B_{avg} indicate a disk spending more time to access a block than other disks. For disk drives, block access time consists of seeking time, rotational delay and transfer time. Among them, the seeking time dominates most of the access time and is easily affected by I/O request pattern (random and sequential). Successive random requests lead to a continuous track seeking and spend plenty of seeking time. Although the random performance of disk drives has been improved by adjacent I/O request merging and elevator algorithm, the seeking operation, determined by application demand, has always existed to lead to a long time transfer of a bit of data in data centers. Therefore, we speculate the I/O workload pattern is the reason of the 'abnormal disk'.

To prove our speculation, we use ioMeter [15] to test the impact of I/O workload pattern on B_{avg} and D_{avg}. In our experiment, we turn up the random/sequential ratio from 0 to 1. We find significant D_{avg} increasing (61.32% to 100%) and B_{avg} decreasing (26493 KB/s to 412 KB/s). Note that ioMeter is a performance benchmark, D_{avg} is always high (greater than 60% in our experiment). But D_{avg} of most drives are much lower in data centers. Even so, we are reasonable to say the I/O workload pattern is one of the most likely reasons of the 'abnormal disk'.

Next, we present the randomness of workload in each interval by the *Intensity of Sequential Request* (ISR) for a convenient statement later. We know disk drives with sequential workload produces a small duty cycle and a large bandwidth in each interval. Thus, we use the ratio of bandwidth (b) to duty cycle (d) to represent the I/O workload pattern.

We explain the feature by the definition of duty cycle and bandwidth. Duty cycle is the period used for data transfer divided by the sampling interval (five minutes in our study). The bandwidth is the bytes of transferred data divided by the interval. Therefore, ISR is a variation of bandwidth, referring to the amount of data divided by the working period instead of the five minutes interval. A small ISR not only means more random requests, but also indicates the 'bandwidth' is low during its working time.

At last, we explain the impact of D_{avg} and the impact of B_{avg} on *AFR* by the randomness of workload. In Figure 2, the D_{avg} of 50% divides drives into a high risk group and a low risk group. That's because a large D_{avg} means more random requests, leading to a frequent mechanism movement of track seeking. For each disk, D_{avg} of 50% is probably the point where random requests begins to show its impact on disk failure. In Figure 7, *AFR* vibrates without an obvious trend when B_{avg} is less than 5000 KB/s. That's because workload contains both random requests and sequential requests when B_{avg} is less than 5000 KB/s. When B_{avg} is greater than 5000 KB/s, the *AFR* decreases since disks with higher B_{avg} serves more sequential requests than random requests. That's the reason of the count-intuitive *AFR* reduction in Figure 7.

Finding 5: Because D_{avg} and B_{avg} affect AFR from different aspects, each of them cannot provide a complete picture on how I/O workload impacts disk reliability. In order to clarify this complex relationship and clearly understand the whole picture, we need a high-level and comprehensive metric beyond D_{avg} and B_{avg} to represent how likely and how seriously an I/O workload does harm to disk health.

B. EFFECT OF WORKLOAD RANDOMNESS

For each disk drive, we use the *AISR* to quantify the impact of random request, shown in Formula 7, which is easily obtained in data centers. We plot the distribution of *AISR* in Figure 12 and the impact of *AISR* on *AFR* in Figure 13. For the same reason as B_{avg}, we truncate 1.34% drives with *AISR* greater than 400 and add them in the last bin. In Figure 12, most disk drives have a low *AISR*. More than half drives have an *AISR* less than 100. That's because large *AISR* only appears on disk drives executing plenty of sequential requests, such as file server and video server, which is a small group in our study. For the rest servers focusing on connection or computation, *AISR* is small.

$$AISR = \frac{1}{n}\sum_{i=1}^{n}\frac{b_i}{d_i} \tag{7}$$

Figure 13 shows a clear reduction of *AFR*. The Pearson's correlation coefficient of −0.939 indicates a strong negative correlation between *AISR* and *AFR*. To improve the reliability of disk drive, we suggest data centers increase the ISR of disk drives. One way is to make as much use of the merging strategy as possible by extending waiting time of I/O request for latency insensitive jobs. In this way, the operating system can merge more close I/O request together to save the time of seeking operation. The *solid-state drive* (SSD), which has a great performance for random requests, could improve ISR as well. Delivering workload with low ISR to SSD and assigning workload of sequential requests to disk drive can

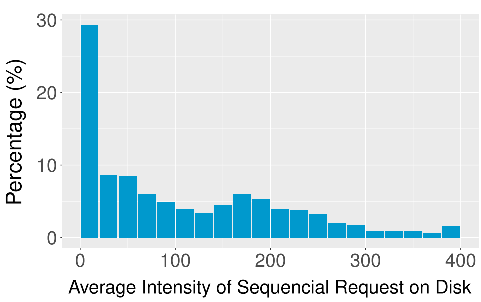

FIGURE 12. Distribution of *Average intensity of sequential request*. Most of disk drives concentrate on a low value.

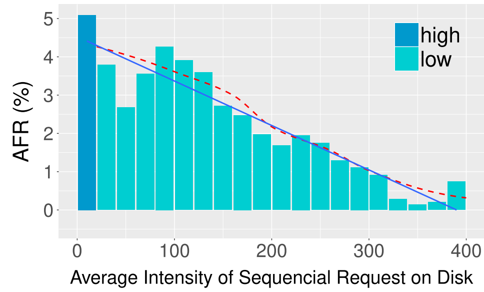

FIGURE 13. ***AFR* vs. *Average intensity of sequential request*. There is a strong positive correlation between *Average intensity of sequential request* and *AFR*. A linear model is used to represent their relationship, shown by the blue line. It fits the *AFR* trend well.**

not only form a heterogeneous storage system to improve storage performance but also be good to storage reliability.

C. DISCUSSION OF WORKLOAD RANDOMNESS

Two metrics (duty cycle and bandwidth) of workload intensity is used to form ISR. However, the forming method of ratio leaves an unknown relationship between ISR and workload intensity since one workload intensity metric divides by another workload intensity metric may mix up the workload intensity, especially when the two intensity metrics have a complex relationship. For example, ISR can remain unchanged with an growing workload intensity where duty cycle grows with bandwidth proportionally. Therefore, we need to study the impact of workload intensity on AISR.

We use D_{avg} and B_{avg} to represent the workload intensity and study their impact respectively. To keep the relative size of the workload intensity, we use the same *AISR* bin as Figure 13 and divide the workload intensity into low, median and high by quantile of 33% and quantile of 66% to represent different level of workload intensity. We plot them in Figure 14 and Figure 15 respectively. From the two figures, we observe workload intensity has an obvious impact on *AFR* at a low *AISR*. When *AISR* is less than 150, the disks with a high-level workload intensity has an greater *AFR* of 7% than the rest disks. While *AISR* is greater than 150,

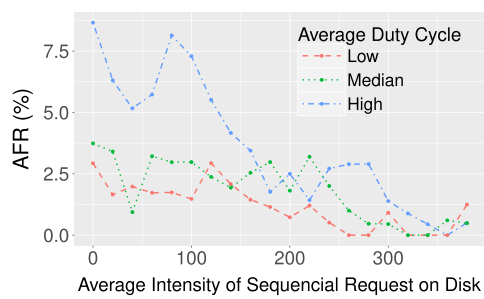

FIGURE 14. ***AFR* vs. *Average intensity of sequential request* under three *Average duty cycle* levels. To avoid significant variability, we use the quantile of *Average duty cycle* (33% and 66%) to divide disk drives. The highest *AFR* occurs with a small *Average intensity of sequential request* and a large *Average duty cycle*.**

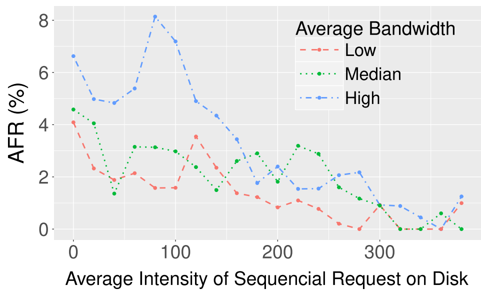

FIGURE 15. ***AFR*** **vs.** ***Average intensity of sequential request*** **under three** ***Average bandwidth*** **levels. To avoid significant variability, we use the quantile of** ***Average bandwidth*** **(33% and 66%) to divide disk drives. The highest** ***AFR*** **occurs with a small** ***Average intensity of sequential request*** **and a large** ***Average bandwidth*****.**

there is no obvious difference between the three lines in both figures. Therefore, disk drives with high workload intensity combined with a small *AISR* are the most vulnerable drives. For these dangerous drivers, we could either reduce their ISR by methods mentioned in last part or alleviate workload intensity by scaling up servers to share jobs. It is interesting to find the impacts of the two intensity metrics are similar in all three levels. Although they have a complex relationship, their impact on the *AFR* is similar. It indicates *AISR* is indeed a comprehensive metric, containing key factors of disk failure from duty cycle and bandwidth.

Based on the strong linear relationship between *AISR* and *AFR*, we build a linear model to quantify the damage of random requests in Equation 8. We estimate the parameters α and β.

$$\alpha * AISR + \beta = AFR \tag{8}$$

With a linear regression technology, we fit the model by α of 0.0116 and β of 4.65 with high confidence. It indicates that *AFR* grows one percent when *AISR* decreases 100. We also plot the line representing the model in Figure 13. It matches well with the actual *AFR* trend. In this model, the maximum *AFR* is 4.65% because of the negative α. In fact, we observe *AFR* greater than 5% in lots of figures above. According to the relationship between *AISR* and workload intensity, we speculate drives with high *AFR* concentrate on the small *AISR* bin with large workload intensity.

As we know, the seeking operation is one of the reasons to explain the relationship between the duty cycle and the bandwidth, although it plays well in the evaluation of *AFR*. A lot of factors can affect the ISR, such as the logical block retries or remapping on failing drives (which will intuitively correlate with failures in the near term regardless of workload), disk cache hits, access sizes, and inner/outer track positioning, etc. Generally, we explain the relationship between workload and *AFR* with one of the possible reasons, and the rest will be studied in our future work.

Finding 6: We propose AISR to integrate the AFR impact from D_{avg} and B_{avg}. With AISR, we can accurately estimate the impact of workload on AFR. It provides a new perspective to manage workload in data centers.

V. RELATED WORK

With the explosive growth of data, how to protect the massive volume of data from disk failures has become a prime interest in recent works. Failure tolerance [16]–[19] is by far the mainstream solution to retain data availability in case of disk failures in data centers. Alternatively, the proactive failure prediction [12], [20]–[23] aims to eliminate the impact before the occurrence of disk failures. While orthogonal to these researches, our work focuses on failure analysis, which plays an important role when making efficient strategies of failure tolerance or building accurate failure prediction model.

A large body of studies try to understand the disk failure from the inherent characteristics [3], [4], [24] and the environment factors [1], [2], [5], [6], [9], [10], [25], [26]. Anderson *et al.* [4] pay attention to the interface of disk drives. Sankar *et al.* [2] consider the influence of temperature at different location granularities. In [5], tradeoffs between energy consumption, environmental conditions, datacenter costs and component reliability are quantified. Few studies focus on the disk workload. To the best of our knowledge, this is the first comprehensive study of uncovering the relationship between the workload and disk failures.

There has been relatively limited work on how I/O workload can influence disk reliability. Yang and Sun [3] and Cole [24], conducted from Quantum and Seagate, attempt to make long-term reliability predictions based on accelerated life tests of small populations. They conclude that power-on-hours, duty cycle and temperature are identified as the key deployment parameters that impact failure rates, each of them having the potential to double failure rates when going from nominal to extreme values. One of the most closely related work [6], from the view of customer Google, concluded that temperature and activity levels of disk drive have less correlation to disk failures. With the conflicting results from the typical works, we try to explore the relationship between the disk I/O workload and the disk reliability. As a case study of user experience, we find both the time of data transfer (duty cycle, used in [24]) and the amount of data transferred (bandwidth, used in [6]) show a relevance to disk failures. Furthermore, we propose a comprehensive metric to summarize individual contributions to the disk failure from them and present different strategies to improve the storage reliability for data centers.

VI. CONCLUSION

In this study, by collecting and investigating about 4 billion drive hours I/O trace of over 500,000 disks in production data centers of Tencent, we report on the characteristics of I/O workload and exploit the relation between I/O workload and disk failure. Our goal is to reveal how I/O workload impacts disk health in data centers. A number of interesting insights into disk reliability in the field are presented.

One of our key findings is that D_{avg} is more relevant to disk failure than B_{avg}. At first, we find D_{avg} shows a positive relationship with AFR. If disk has an D_{avg} greater than 50%, its AFR is 3.47 times as much as AFR of disk with D_{avg} less than 50%. It implies that, different from CPU resource management, roughly increasing disk utilization (i.e., D_{avg}) is probably not a suitable way to reduce TCO of data centers. There is a trade-off between benefit of high disk utilization and expensive recovery cost of failed disks. Then, our study shows that B_{avg} might not be a good indicator of disk failure. It displays a weak impact on AFR when it is less than 5000 KB/s and shows a negative relationship with AFR when B_{avg} is greater than 5000 KB/s. Since D_{avg} and B_{avg} are both indicators of workload intensity, the widely divergent impact from them leads us to study their relationship.

Another key findings is that the expected linear relationship between D_{avg} and B_{avg} only exists when D_{avg} is less than 50%. When D_{avg} is greater than 50% , B_{avg} starts to decrease. Based on understanding of seeking tracks, we attribute the dissymmetry relationship between D_{avg} and B_{avg} to the randomness of I/O workload. We thus propose a new metric $AISR$ to describe the randomness of I/O workload and verify its strong correlation with AFR. As an easily collected metric in data centers and a key characteristic of I/O workload from the view point of disk reliability, $AISR$ can help to provide a new perspective of I/O scheduling in data centers.

Although this work is an exploratory step towards improving disk reliability via I/O workload, we believe insights in our study will spur research community to carefully consider the assignment of disk I/O workload. In future work, we will consider more metrics of workload to mine valuable findings and implement them in data centers.

ACKNOWLEDGEMENTS
The authors are grateful to Xiaosong Ma for her help and insights. The authors also thank group members of system monitoring in Tencent who help them with data.

REFERENCES
[1] B. Schroeder and G. A. Gibson, "Disk failures in the real world: What does an MTTF of 1 000 000 hours mean to you," in *Proc. USENIX Conf. File Storage Technol.*, 2007, pp. 1–16.
[2] S. Sankar, M. Shaw, K. Vaid, and S. Gurumurthi, "Datacenter scale evaluation of the impact of temperature on hard disk drive failures," *ACM Trans. Storage*, vol. 9, no. 2, p. 6, 2013.
[3] J. Yang and F.-B. Sun, "A comprehensive review of hard-disk drive reliability," in *Proc. Annu. Rel. Maintainab. Symp.*, Jan. 1999, pp. 403–409.
[4] D. Anderson, J. Dykes, and E. Riedel, "More than an interface-SCSI vs. ATA," in *Proc. USENIX Conf. File Storage Technol.*, 2003, pp. 245–257.
[5] I. Manousakis, S. Sankar, G. McKnight, T. D. Nguyen, and R. Bianchini, "Environmental conditions and disk reliability in free-cooled datacenters," in *Proc. USENIX Conf. File Storage Technol.*, 2016, pp. 53–65.
[6] E. Pinheiro, W. D. Weber, and L. A. Barroso, "Failure trends in a large disk drive population," in *Proc. USENIX Conf. File Storage Technol.*, 2007, pp. 17–23.
[7] (2017). *Number of Monthly Active WeChat Users From 2nd Quarter 2010 to 2nd Quarter 2017 (in Millions)*. [Online]. Available: https://www.statista.com/statistics/255778/number-of-active-wechat-messenger-accounts/
[8] R. Faulkner and R. Gomes, "The process file system and process model in UNIX system V," in *Proc. USENIX Winter Conf.*, 1991, pp. 243–252.
[9] W. Jiang, C. Hu, Y. Zhou, and A. Kanevsky, "Are disks the dominant contributor for storage failures: A comprehensive study of storage subsystem failure characteristics," in *Proc. USENIX Conf. File Storage Technol.*, 2008, pp. 111–125.
[10] N. El-Sayed, I. A. Stefanovici, G. Amvrosiadis, A. A. Hwang, and B. Schroeder, "Temperature management in data centers: Why some (might) like it hot," *ACM SIGMETRICS Perform. Eval. Rev.*, vol. 40, no. 1, pp. 163–174, 2012.
[11] Y. Kim, S. Gurumurthi, and A. Sivasubramaniam, "Understanding the performance-temperature interactions in disk I/O of server workloads," in *Proc. Int. Symp. High Perform. Comput. Archit.*, Feb. 2006, pp. 176–186.
[12] J. Li *et al.*, "Hard drive failure prediction using classification and regression trees," in *Proc. Annu. IEEE/IFIP Int. Conf. Dependable Syst. Netw.*, Jun. 2014, pp. 383–394.
[13] B. Zhu, G. Wang, X. Liu, D. Hu, S. Lin, and J. Ma, "Proactive drive failure prediction for large scale storage systems," in *Proc. IEEE Symp. Mass Storage Syst. Technol.*, May 2013, pp. 1–5.
[14] J. M. Kaplan, W. Forrest, and N. Kindler, "Revolutionizing data center energy efficiency," McKinsey & Company, New York, NY, USA, Tech. Rep., 2008.
[15] (1997). *Iometer is an I/O Subsystem Measurement and Characterization Tool*. [Online]. Available: http://www.iometer.org
[16] A. Ma *et al.*, "RAIDShield: Characterizing, monitoring, and proactively protecting against disk failures," in *Proc. USENIX Conf. File Storage Technol.*, 2015, pp. 241–256.
[17] M. Li and P. P. C. Lee, "Toward I/O-efficient protection against silent data corruptions in RAID arrays," in *Proc. IEEE Symp. Mass Storage Syst. Technol.*, Jun. 2014, pp. 1–12.
[18] T.-Y. Youn and K.-Y. Chang, "Bi-directional and concurrent proof of ownership for stronger storage services with de-duplication," *Sci. China Inf. Sci.*, vol. 61, no. 3, p. 032107, 2018.
[19] Z. Wang, Z. Han, and J. Liu, "Public verifiability for shared data in cloud storage with a defense against collusion attacks," *Sci. China Inf. Sci.*, vol. 59, no. 3, p. 039101, 2016.
[20] Y. Wang, E. W. M. Ma, T. W. S. Chow, and K.-L. Tsui, "A two-step parametric method for failure prediction in hard disk drives," *IEEE Trans Ind. Informat.*, vol. 10, no. 1, pp. 419–430, Feb. 2014.
[21] J. S. Plank and M. Blaum, "Sector-disk (SD) erasure codes for mixed failure modes in RAID systems," *ACM Trans. Storage*, vol. 10, no. 1, pp. 4–14, 2014.
[22] X. Luo and J. Shu, "Load-balanced recovery schemes for single-disk failure in storage systems with any erasure code," in *Proc. IEEE Int. Conf. Parallel Process.*, Oct. 2013, pp. 552–561.
[23] J. Xiao, Z. Xiong, S. Wu, Y. Yi, H. Jin, and K. Hu, "Disk failure prediction in data centers via online learning," in *Proc. IEEE Int. Conf. Parallel Process.*, Aug. 2018, pp. 1–10.
[24] G. Cole, "Estimating drive reliability in desktop computers and consumer electronics systems," Seagate Technol. PLC, Cupertino, CA, USA, Seagate Technol. Paper TP-338.1, 2000.
[25] J. G. Elerath and S. Shah, "Server class disk drives: How reliable are they?" in *Proc. Annu. Rel. Maintainab. Symp.*, Jan. 2004, pp. 151–156.
[26] R. Birke, I. Giurgiu, L. Y. Chen, D. Wiesmann, and T. Engbersen, "Failure analysis of virtual and physical machines: Patterns, causes and characteristics," in *Proc. Annu. IEEE/IFIP Int. Conf. Dependable Syst. Netw.*, Jun. 2014, pp. 1–12.

SONG WU received the Ph.D. degree in computer science from the Huazhong University of Science and Technology in 2003. He is currently a Professor with the Computer School, Huazhong University of Science and Technology, where he is also with the members of the System Group, SCTS/CGCL. He is the Director of the Parallel and Distributed Computing Institute. His research interests include cloud computing, system virtualization, datacenter management, storage system, and in-memory computing.

YUSHENG YI received the B.S. degree from the School of Mathematics and Statistics, Huazhong University of Science and Technology, Wuhan, China, in 2011, where he is currently pursuing the Ph.D. degree with the School of Computer Science and Technology. His current research interests include disk reliability, failure prediction, and data mining.

JIANG XIAO is currently an Associate Professor with the School of Computer Science and Technology, Huazhong University of Science and Technology, Wuhan, China. She currently involved in research on wireless indoor localization, smart sensing, big data analysis, and distributed computing.

HAI JIN received the Ph.D. degree in computer engineering from the Huazhong University of Science and Technology (HUST), China, in 1994. Between 1998 and 2000, he was with The University of Hong Kong. Between 1999 and 2000, he was a Visiting Scholar with the University of Southern California. He is currently a Cheung Kung Scholars' Chair Professor of computer science and engineering with HUST. He has co-authored 15 books and authored over 700 research papers. His research interests include computer architecture, virtualization technology, cluster computing and cloud computing, peer-to-peer computing, network storage, and network security. He is a senior member of IEEE. He was a recipient of the German Academic Exchange Service Fellowship to visit the Technical University of Chemnitz, Germany, in 1996. He received the Excellent Youth Award from the National Science Foundation of China in 2001. He is the Chief Scientist of ChinaGrid, the largest grid computing project in China, the National 973 Basic Research Program Project of Virtualization Technology of Computing System, and Cloud Security.

MAO YE received the B.S. and M.S. degrees from the School of Computer Science and Technology, Huazhong University of Science and Technology, Wuhan, China, in 2012 and 2014, respectively. She is currently with Tencent Inc., Shenzhen, China.