An Efficient Penalty-Aware Cache to Improve the Performance of Parity-based Disk Arrays under Faulty Conditions

Shenggang Wan, Xubin He, *Senior Member, IEEE*, Jianzhong Huang, Qiang Cao, *Member, IEEE,* Shiyi Li, and Changsheng Xie, *Member, IEEE*

Abstract—The buffer cache plays an essential role in smoothing the gap between the upper-level computational components and the lower-level storage devices. A good buffer cache management scheme should be beneficial to not only the computational components, but also the storage components by reducing disk I/Os. Existing cache replacement algorithms are well optimized for disks in normal mode, but inefficient under faulty scenarios, such as a parity-based disk array with faulty disk(s).

To address this issue, we propose a novel penalty-aware buffer cache replacement strategy, named Victim Disk(s) First (VDF) cache, to improve the reliability and performance of a storage system consisting of a buffer cache and disk arrays. VDF cache gives higher priority to cache the blocks on the faulty disks when the disk array fails, thus reducing the I/Os addressed directly to the faulty disks.

To verify the effectiveness of the VDF cache, we have integrated VDF into the popular cache algorithms LFU and LRU, named VDF-LFU and VDF-LRU, respectively. We have conducted intensive simulations as well as a prototype implementation for disk arrays to tolerate one disk failure (RAID-5) and two disk failures (RAID-6). The simulation results have shown that VDF-LFU can reduce disk I/Os to surviving disks by up to 42.3% in RAID-5 and 50.7% in RAID-6, and VDF-LRU can reduce those by up to 36.2% in RAID-5 and 48.9% in RAID-6. Our measurement results also show that VDF-LFU can speed up the online recovery by up to 46.3% in RAID-5 and 47.2% in RAID-6 under spare-rebuilding mode, or improve the maximum system service rate by up to 47.7% in RAID-5 under degraded mode without a reconstruction workload. Similarly, VDF-LRU can speed up the online recovery by up to 34.6% in RAID-5 and 38.2% in RAID-6, or improve the system service rate by up to 28.4% in RAID-5.

Index Terms—Buffer Cache Replacement Algorithm, Reconstruction, Reliability, Performance.

✦

1 INTRODUCTION

To reduce the number of I/O requests to the low level storage device, such as disk arrays, a cache is widely used and many cache algorithms exist to hide the long disk latencies. These cache algorithms work well for disk arrays under normal fault-free mode. However, when some disks in a disk array fail, the RAID may still work under this faulty scenario, either a *spare-rebuilding mode* with online reconstruction or a *degraded mode* without online reconstruction. The cost of a miss to faulty disks might be dramatically different compared to the cost of a miss to surviving disks. Existing cache algorithms cannot capture this difference because they treat the underlying (faulty or surviving) disks in an array the same.

We take an example as shown in Figure 1, which illustrates two different cache miss situations in a storage subsystem composed of a parity-based RAID with one faulty disk in degraded mode. As shown in Figure 1(a), the missed data resides in the faulty disk. The RAID controller accesses the surviving disks to fetch all data and parity in the same stripe to regenerate the lost data. Therefore, to service one cache miss, several read requests are needed depending on the RAID organization. However, if the missed data is in a surviving disk as shown in Figure 1(b), only one read request to the corresponding surviving disk is generated. Similar situations are observed in spare-rebuilding mode. A simple analysis shows that in a RAID-5 system consisting of n disks, when a disk fails, the cost to fetch data from a faulty disk might be $n-1$ times higher than the cost to access data from a surviving disk. This extra disk I/O activity will reduce the effective array bandwidth available for reconstruction or user access.

When a disk array starts online reconstruction, it uses up regular bandwidth. Compared to offline reconstruction, during the process of online reconstruction, the user workflow interferes with the reconstruction workflow. As a result, the online reconstruction duration grows significantly compared to offline reconstruction. Wu et al. [1] point out that, in a heavy user workflow, the duration of online reconstruction would grow as much as 70 times as that of the offline reconstruction. In this case, more requests to the surviving disks, caused by user requests, reduce the available reconstruction bandwidth and lengthen the reconstruction duration, which reduces the reliability of the storage system.

On the other hand, in a degraded mode without a reconstruction workflow, a miss to faulty disks would cause all the surviving data in the same parity chain (stripe in RAID-5) to be read and add additional workflow to surviving disks. With a decreasing serviceability and an increasing user workflow caused by misses to faulty disks, the storage subsystem might be overloaded under a heavy user workflow.

Furthermore, considering the critical cases in disk arrays tolerating two or more concurrent failures, such as RAID-6, with the increased number of faulty disks, the misses of

- S. Wan, J. Huang*, Q. Cao, S. Li, and C. Xie are with the School of Computer Science and Technology, Huazhong university of science and technology, Wuhan,Hubei 430074, China. E-mail: swan3@vcu.edu, *corresponding author: hjzh@hust.edu.cn, caoqiang@hust.edu.cn, lsy011025@gmail.com, cs_xie@hust.edu.cn.

- X. He is with the Department of Electrical and Computer Engineering, Virginia Commonwealth University, 601 West Main St.,Richmond, VA 23284, USA. Email: xhe2@vcu.edu. S. Wan is currently a visiting scholar at VCU working with X. He.

(a) A read miss to a faulty disk might result in several additional read requests to the surviving disks.

(b) A read miss to a surviving disk would result in only one request to the corresponding surviving disk.
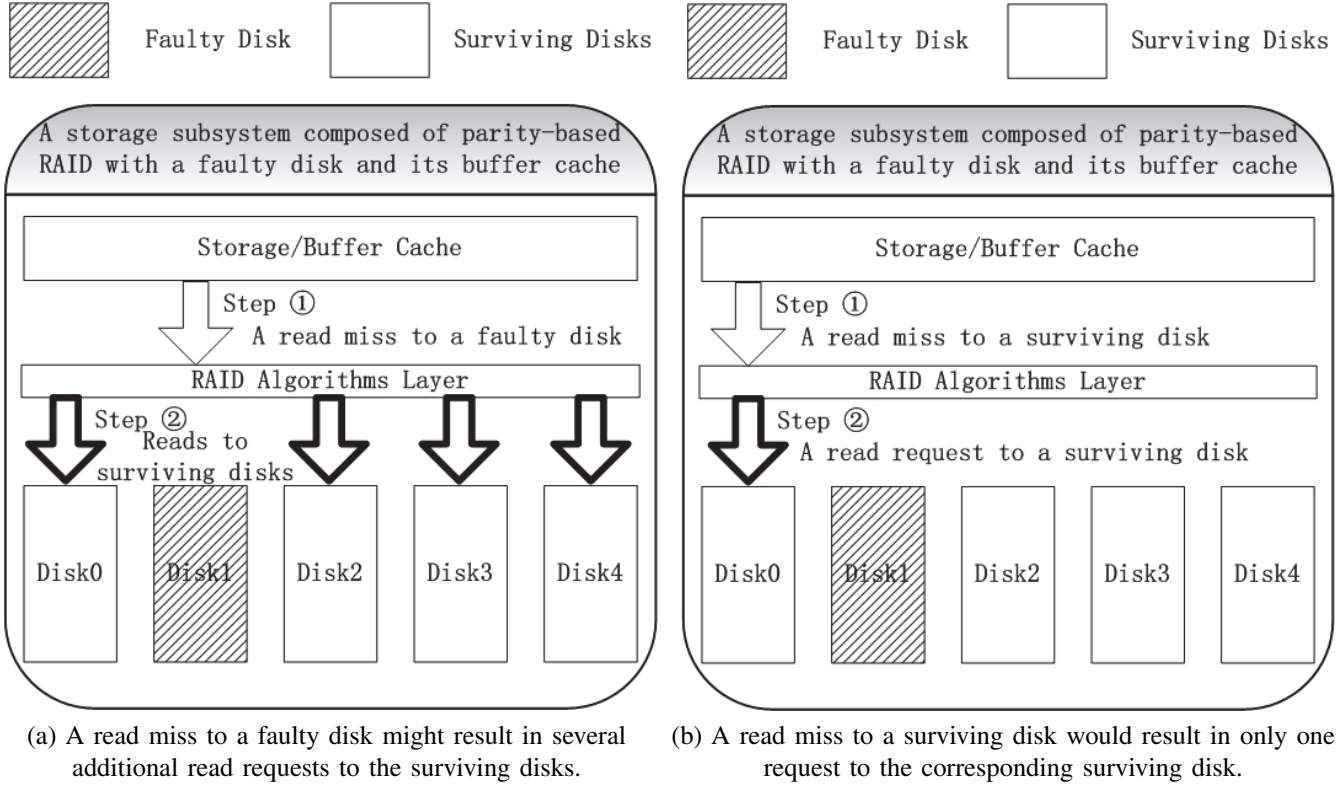
Fig. 1. Two typical cache-miss situations in a storage subsystem composed of a parity-based RAID in degraded mode.

requests to faulty disks grow accordingly. Thus, the disk arrays would suffer from much more additional requests than that under the case of single disk failure. As a result, the reliability and the performance problems would be more significant.

Therefore, in parity-based disk arrays under faulty conditions, a miss to faulty disks is much more expensive than a miss to surviving disks. Based on this observation, we propose a penalty-aware buffer cache replacement strategy, named Victim (or faulty) Disk(s) First cache, or VDF for short, to improve the performance of storage subsystem composed of a parity-based disk array and its buffer cache. The basic idea is to design a cache scheme to treat the faulty disks more favorably, or give higher priority to cache the data associated with the faulty disks. The goal of this scheme is to reduce the cache miss directed to the faulty disk, and thus to reduce the I/O requests to the surviving disks overall. Reduced disk I/O caused by the user workflow will (1) improve the performance of the disk array, and (2) allow more bandwidth for online reconstruction which in turn speeds up the recovery, and thus improves the reliability.

We make the following four contributions in this paper:

1) We propose a new metric, Requests Generation Ratio or RGR, to capture the disk I/O activities of user workflows on the surviving disks when a storage system is under faulty conditions. This would directly influence the maximum bandwidth for reconstruction in spare-rebuilding mode and the bandwidth available to user workflows in degraded mode.

2) We develop a novel cache replacement scheme, VDF, by

giving higher priority to cache the data associated with the faulty disks, to minimize the RGR. VDF is flexible, and could be integrated into existing cache algorithms such as LRU and LFU.

3) We conduct intensive simulation to verify the effectiveness of VDF under different workloads and different RAID levels. The simulation results show that VDF-LRU can reduce overall disk I/Os to surviving disks by up to 36.2% in RAID-5 and 48.9% in RAID-6. VDF-LFU can reduce those by up to 42.3% in RAID-5 and 50.7% in RAID-6.

4) We build a prototype of VDF to evaluate its effectiveness in the Linux software RAID system. As a result, VDF-LFU can speed up the online recovery by up to 46.3% in RAID-5 and 47.2% in RAID-6 under spare-rebuilding mode, or improve the maximum system service rate by up to 47.7% in RAID-5 under degraded mode. Similarly, VDF-LRU can speed up the online recovery by up to 34.6% in RAID-5 and 38.2% in RAID-6, or improve the system service rate by up to 28.4% in RAID-5.

The rest of the paper is organized as follows: Section 2 gives a brief overview of the background and related work. In Section 3, we describe our new metric, RGR, and the design of VDF. Detailed case studies of VDF cache for disk arrays tolerating single and concurrent failures are given in Section 4; we describe integrating VDF into two typical cache-replacement algorithms, LRU and LFU. We provide our simulation results in Section 5, and prototyping and measurement results in Section 6. We conclude our paper in Section 7.

2 BACKGROUND AND RELATED WORK

In this section we briefly overview some background materials and related work.

2.1 Optimizations of Disk Arrays under Faulty Conditions

Redundant Arrays of Independent Disks RAID [2] are popular solutions to provide high performance and reliability for today's storage systems. Depending on its organizations, RAID could prevent data loss incurred by disk failures and even offer online services under faulty conditions. With faulty disk(s), these RAIDs would work in a spare-rebuilding mode to support online reconstruction, or in a degraded mode without reconstruction.

RAID can offer continuous online services even in faulty mode. However, the recovery workload and user request can interfere with each other, and lead to longer recovery times. Many solutions are proposed to solve this problem, such as optimizations of data/parity/spare layout [3]–[7], reconstruction workload [8]–[12], and user workload [1], [13], [14].

Menon et al. present a method to distribute spares to all disks, which would not only reduce the lost data per disk but also parallelize the reconstruction [4]. Holland et al. [3] propose a trade-off between RAID-1 (mirror) and RAID-5, named parity declustering, to balance the storage efficiency and the recovery performance. Xin et al. use a RUSH-like hash algorithm to evenly distribute data, parity, and spares among the nodes in a distributed environment [5].

The track-based recovery (TBR) [8] algorithm provides a trade-off between block-based recovery and cylinder-based recovery, and balances the user response time and the recovery duration. However, TBR requires much more buffer space compared to block-based recovery. The pipelined recovery (PR) scheme [9] addresses this problem, and significantly reduces the buffer requirements close to that of the block-based recovery algorithms. The disk-oriented recovery (DOR) algorithm [10] rebuilds the array at the disk-level instead of the stripe-level. With this approach, DOR could absorb the bandwidth of the array as much as possible. The popularity-based recovery (PRO) algorithm [11], [12], builds upon the DOR algorithm, further improving the recovery performance by utilizing the spatial locality of user requests.

Two techniques named *redirection of reads* and *piggybacking of writes* [13] are proposed to reduce the user workflow by employing the reconstructed spare disk to absorb parts of the requests to the faulty disk. However, they need to maintain a bitmap in the dedicated cache in the RAID device to record the reconstruction status; as the increasing of disk size, a fine-granularity bitmap would consume too much memory, and increase synchronization costs. For example, a bitmap with granularity of 4KB for a 2TB disk would require 64MB of memory, which limits the use of piggybacking of writes. During the reconstruction, at most a coarse-granularity bitmap could be used only to redirect reads. MICRO [14] achieves improved recovery performance by writing back the in-memory surviving data of the faulty disks into a spare disk first and using a file popularity table to find the hotspot. MICRO treats

all the blocks in the cache equally, which is similar to the general cache-replacement algorithms and has the same limitations. WorkOut [1], an array-cache-array method, offloads the write requests and popular read requests to another disk array. As a result, WorkOut speeds up the recovery process and improves the user response time. However, WorkOut requires another disk array to help with the reconstruction and need maintain an addressing translation map, which might be much larger than a fine-granularity bitmap, in the dedicated cache on the RAID device. This suffers from the same problem as redirection of reads and piggybacking of writes.

2.2 Buffer Cache Replacement Algorithms

RAID-based storage systems usually work together with the buffer cache. To improve the efficiency of the buffer cache, researchers have proposed many cache-replacement algorithms, such as LRU [15], LFU, FBR [16], LRU-k [17], [18], 2Q [19], LRFU [20], [21], MQ [22], [23], LIRS [24], [25], ARC [26], DULO [27], DISKSEEN [28] and more. Each cache-replacement algorithm weigh the cached blocks with a different method, such as access interval, access frequency and so on, then decide which to evict.

The LRU (Least-Recently-Used) algorithm is one of the most popular and effective policies for buffer cache management. When a block needs to be inserted into the cache, the candidate to be evicted is the block which is least recently used. That is to say, the weight of the cached blocks in LRU is its last access timestamp. The block with the smallest last access timestamp is evicted. The LFU (Least-Frequently-Used) algorithm replaces the least frequently used block. In other words, the weight of the cached blocks in LFU is its number of accesses. The block with the smallest number of accesses is evicted. Other algorithms, such as LRU-k, 2Q, LRFU, MQ, LIRS, and ARC, integrate LRU and LFU algorithm together and demonstrate good performance under various scenarios. DULO and DISKSEEN consider both temporal and spatial locality when a block need to be replaced.

However, the above cache-replacement algorithms work well when the RAID system is under normal operating mode. When some disks in the RAID system fail, it runs under faulty condition, but the buffer cache layer is not aware of the underlying failures in RAID and thus the existing cache algorithms do not work well as explained in Section 1. This motivates us to propose VDF: a cache scheme to treat the faulty disks more favorably, or give a higher priority to cache the data associated with faulty disks. The goal is to reduce the cache misses directed to the faulty disk and thus to reduce the I/O requests to the surviving disks overall. As our VDF only increases the weight of blocks in the faulty disks, theoretically it could work with the above-mentioned general cache replacement algorithms.

2.3 RAID-6 Codes

According to the Storage Networking Industry Association (SNIA), A RAID-6 is "any form of RAID that can continue to execute read and write requests to all of an array's virtual disks in the presence of two concurrent disk failures" [29].

TABLE 1
Variables and Definitions

Symbols	Definition
C	Total number of blocks in the buffer cache
T	Total number of data blocks in a disk array
B_i	Data block i
p_i	Access probability of each block B_i
MP_i	Miss penalty of each block B_i
BW	Total serviceability of all surviving disks in terms of I/O bandwidth
BW_U	I/O bandwidth available to user workload, or service rate of the system from the user's point of view
BW_R	I/O bandwidth for an online reconstruction workload
RGR	The ratio of the # of requested blocks to surviving disks and the # of requested blocks to buffer cache
Q	Total amount of data from surviving disks to reconstruct faulty disks

Compared to the single disk failure tolerated disk arrays, RAID-6 disk arrays are more reliable and gaining increased popularity, particularly in large scale data centers due to the increased possibility of multiple failures in such systems.

RAID-6 implementations can be based on various codes, including Reed-Solomon codes [30], Cauchy Reed-Solomon codes [31], EVENODD codes [32], RDP codes [33], Blaum-Roth codes [34], Liberation codes [35], Cyclic codes [36], X-codes [37], P-codes [38], DFDP [39], HoVer codes [40], Code-M [7], H-Code [41], and HDP-Code [42].

Reed-Solomon codes are based on addition and multiply operations over certain finite fields GF(2^w). The addition operation over GF(2^w) can be implemented by exclusive-OR operation. However, the multiply operation should be implemented using tables with much more complex. Cauchy Reed-Solomon codes address this problem and improve Reed-Solomon codes by changing the complex multiply operation into additional XOR operations.

EVENODD is a special erasure coding technology only for the case of RAID-6. It contains two types of parity: the P parity, which is just like the horizontal parity in RAID-4, and the Q parity, which is generated by the elements on the diagonal. RDP is another special erasure coding technology only for the case of RAID-6. The P parity of RDP is the same as EVENODD, but a different construction of the Q parity brings the optimal construction and reconstruction computational complexity.

A special class of erasure coding technologies called *Minimum Density* (MD) codes are proposed to achieve optimal short write/update performance. Blaum *et al* [34] point out in a typical horizontal code defined by an m-row-k-column matrix of data elements for RAID-6, if the P parity is fixed to be horizontal parity, then there must be at least $mk + k - 1$ elements joining in the generation of the Q parity in order to make the code be lowest density. Blaum-Roth, Liberation, and Liber8tion codes are all minimum density. Compared with other horizontal codes for RAID-6, minimum density codes share a common advantage that they have the near-optimal short write/update complexity.

X-codes, Cyclic codes, and P-codes are vertical codes, for their parities are not in the individual redundant disk drives, but dispersed over all the disk drives. This fact gives the two codes optimal computational and update complexity [38]. H-Code and HDP-Code are hybrid codes, taking advantages of both horizontal codes and vertical codes. HoVer codes and Code-M are non-MDS codes to tradeoff the spatial utilization for reconstruction/decoding cost.

Although Reed-Solomon codes have some significant flaws, they are still widely used in many systems, such as the RAID-6 implementation in Linux kernel. And our evaluations on RAID-6 are based on the RS codes.

3 DESIGN OF VDF

In this section, we propose a new metric to describe the cache efficiency of disk I/O activities. We show how to use it to evaluate disk arrays under faulty conditions, and then we describe our VDF scheme. Before our discussion, we summarize the symbols in Table 1.

3.1 RGR: A New Metric to Evaluate Cache Performance with Various Miss Penalty

Traditional cache replacement algorithms are essentially evaluations on blocks access probability, which is typically the frequency of accesses to certain blocks in I/O workloads [17], [18], [24], [25], assuming that the penalty of each miss at the same level is the same. However, in parity-based RAID with faulty disk(s), the penalty of a miss to the lost data in the faulty disks might be much more expensive than that of a miss to surviving data. Therefore, from the aspect of a RAID device, the buffer cache performance should not be simply evaluated by the traditional metrics such as Hit Ratio or Miss Ratio, particularly when the RAID is under faulty conditions. To address this issue, we propose a new metric called *Requests Generation Ratio* or RGR to evaluate the cache performance from the view point of a faulty RAID device. RGR is defined in Equation 1, It represents the disk activities to service an I/O request to the buffer cache. Ideally, if all I/O requests are serviced by the buffer cache, RGR will be 0 (no disk I/Os are generated). For missed I/O requests, RGR will be different depending on the penalty to each underlying disk. For example, in Figure 1(a) the RGR of a miss to the faulty disk is 4 because 4 disk I/Os are generated to service the missed I/O request, and in Figure 1(b), the RGR of a miss to surviving disks is 1.

$$RGR = \frac{\#\ of\ block\ requests\ to\ surviving\ disks}{\#\ of\ block\ requests\ to\ buffer\ cache} \quad (1)$$

To calculate the RGR, we assume a parity-based RAID of T data blocks with a buffer cache of C blocks. The access probability of a block B_i is p_i, where $0 \leq i \leq T - 1$, with a miss penalty of MP_i in terms of the total requested blocks to surviving disks caused by a miss. From the viewpoint of a certain workload, p_i actually represents the ratio of the number of request on B_i and the number of total block requests. If block B_i is not referenced in this workload, p_i should be 0. As we have mentioned above, different cache algorithms evaluate p_i with different approaches in runtime environments. If a block is serviced by the cache, the corresponding miss penalty $MP_i = 0$. Therefore, the RGR of the next block request can be described by the following Equation 2.

$$RGR = \sum_{i=0}^{T-1} (p_i \times MP_i) \qquad (2)$$

3.2 Use RGR to Evaluate the Cache Efficiency in Faulty Mode

Consider a system composed of a buffer cache and a RAID in faulty mode which service a certain user workload. We have the following symbols. First, the total serviceability of all surviving disks is BW in terms of I/O bandwidth. Second, the unfiltered user workload would take BW_U bandwidth, which is the service rate of the system from a user's perspective. The average RGR of the buffer cache is RGR. Therefore, the filtered user workload should take about $BW_U \times RGR$ bandwidth. Third, all the remaining bandwidth BW_R of all surviving disks could be utilized for reconstruction. Lastly, the total amount of surviving data for reconstruction is Q. Equation 3 describes the relationships among BW, BW_U, RGR, and BW_R.

$$BW = BW_U \times RGR + BW_R \qquad (3)$$

We first consider the spare-rebuilding mode. The surviving disks would suffer from more requests as explained in Section 1. It means the I/O bandwidth available for reconstruction on the surviving disks would be less than the I/O bandwidth for reconstruction on the spare disk. The total amount of requested data for reconstruction on each disk (including the surviving disks and the spare disks) is the same. Therefore, to the online recovery process, the I/O bandwidth for reconstruction on the surviving disks is the bottleneck. The reconstruction duration RD could be described with Equation 4.

$$RD = \frac{Q}{BW - BW_U \times RGR} \qquad (4)$$

From Equation 4, we can find that, if Q, BW, and BW_U are fixed, with the decreasing RGR, the reconstruction duration RD (and thus MTTR) decreases. Therefore, to minimize the MTTR, we should minimize the RGR.

We next consider the degraded mode without reconstruction. Each surviving disk would suffer from the extra requests caused by the access to faulty disks. The filtered user workload should not exceed the total serviceability of all surviving disks. In another words, the maximum unfiltered user workload

BW_U should not exceed $\frac{BW}{RGR}$. Therefore, we should minimize the RGR to maximize the system serviceability, which is described with a maximum BW_U.

From the above discussion, we notice that compared to the traditional metrics on cache evaluation, such as miss ratio, RGR is useful to demonstrate two important indicators of a faulty disk array more clearly and directly. One is the reconstruction time which is directly related to MTTR and affects the system reliability. The other is the throughput that indicates the performance of the storage system.

3.3 VDF Cache

Based on above analysis, we propose our VDF cache aiming at reducing the RGR for parity-based RAID under faulty conditions, either to enhance the system reliability by speeding up the reconstruction process in spare-rebuilding mode, or to improve the system performance by increasing the system serviceability in degraded mode without reconstruction workloads. As it works in the buffer cache level, VDF is practical and will not suffer from the problems of the small dedicated cache in a RAID controller.

Cache replacement algorithms are essentially evaluations on access probability of cached blocks. Once a miss occurs, typically a block should be evicted from cache, and the missed block would be loaded to the free space. General replacement algorithms evict the block with the smallest access probability to reduce the total access probability of the remaining blocks out of buffer cache. However, to minimize the RGR, the eviction of a block should not only be determined by the access probability but also by the miss penalty. To integrate VDF into any cache replacement algorithms, we take the same evaluation approach of access probability (p_i) for each cached block as the original algorithm. A cache replacement operation indicates a block is evicted from cache and another block is fetched to cache. Since the block fetched to the cache is known (the missing block), the influence on RGR will be determined by which block is evicted. When a block resides in cache, its MP_i is zero. If it is evicted from cache, its MP_i increases depending on where it is (in a faulty or surviving disk). The eviction of a block will increase the total RGR by $(MP_i - 0) \times P_i$. Therefore if we choose a block with minimum $p_i \times MP_i$ to evict, it has the least impact on increasing RGR and thus to keep RGR minimal.

4 CASE STUDIES OF VDF

To verify the effectiveness of VDF, we apply VDF in disk arrays of n disks to tolerate one disk failure (RAID-5) and two concurrent disk failures (RAID-6). We focus on read operations for two reasons: first, in many applications, users are typically sensitive to read latency, particularly in a disk array under faulty conditions; second, in many storage systems, independent non-volatile memory is deployed as a write cache to enhance the reliability and this cache uses a dedicated write cache algorithm. For cache replacement algorithms, we consider the simple while popular LRU and LFU algorithms.
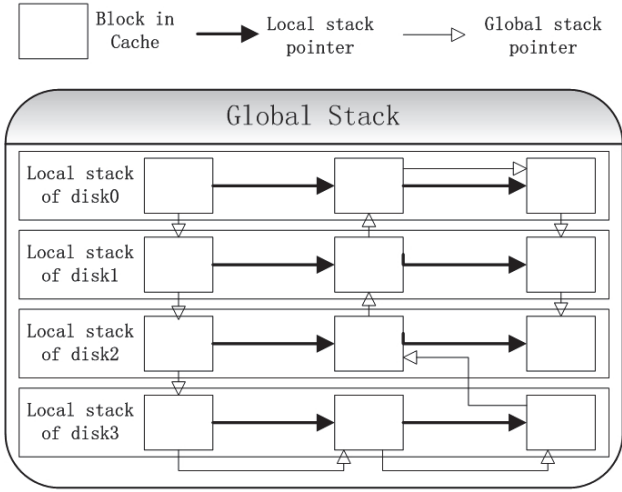
Fig. 2. VDF implementation with two types of stacks.

4.1 Integrate VDF into LRU and LFU

Although VDF cache can cooperate with caches at other levels by adjusting the miss penalty of blocks, for demonstration purposes we just consider a one level buffer cache above the disk array. Therefore, the miss penalty of blocks on the faulty disk would be $n - 1$ for RAID-5 and $n - 2$ for RAID-6 in our following discussion, which means one cache miss to the faulty disk will result in $n - 1$ or $n - 2$ I/O requests to the surviving disks, respectively. To integrate VDF with any cache algorithm, the access probability should be evaluated with a quantitative approach. Different cache replacement algorithms evaluate the access probability of blocks using different approaches. Most existing cache management algorithms can be categorized into LRU-like and LFU-like algorithms. In LRU-like algorithms, the weight of blocks is often evaluated by the access time interval. As it is costly to record the real access timestamp, a straight alternative is to record the access sequence number, and use the reciprocal of the interval access sequence number as the access probability. This approach is widely used in many LRU-like algorithms [17], [18], [20], [21]. In LFU-like algorithms, the weight of blocks is majorly evaluated by the access frequency. Thus, to integrate VDF into these LFU-like algorithms, it needs only to keep the original evaluation approach. Furthermore, different from the access sequence number, the access frequency of two blocks might be the same. Therefore, in VDF based LFU-like algorithms, the access sequence number is also employed for choosing which block to evict with the same access frequency. In VDF cache the access probability of a block would not be the absolute but the relative value, because both the reciprocal of interval of access sequence number and the access frequency are actually the relative values.

The conversion from the original cache algorithms to the VDF-based algorithms should be smooth, because VDF takes effect in faulty mode. In other words, the buffer cache should be managed with the original algorithms in fault-free mode, and the VDF policy becomes effective when disk failures occur. Thus, a smooth runtime conversion between the original algorithm and the VDF-based algorithm is needed, which is

quite different from the general cache algorithms. Therefore, in VDF-based algorithm, two types of stacks are employed to achieve the smooth runtime conversion: one is the global stack (GS) which is similar to the stack in a general algorithm such as global LRU stack, and the other is the local stack (LS) holding the blocks on the same disk in cache. All blocks should be in two types of stacks concurrently as shown in Figure 2. When the system works in fault-free mode, it evicts the block with the smallest access probability (p_i) at the bottom of the GS stack. Once a disk array drops to a faulty mode, it evicts the block with the smallest $p_i \times MP_i$ at the bottom of each LS stack instead of evicting the block at the bottom of the GS stack.

Algorithm 1: VDF-LRU for disk arrays of n disks tolerating one or two failures

---

Input: The request stream $x_1, x_2, x_3, ..., x_i, ...$
VDF_LRU_Replace(x_i){
/*For every $i \geq 1$ and any x_i, one and only one of the following cases must occur.*/
if x_i is in $LS_k, 0 \leq k < n$ **then**
    /*A cache hit has occurred.*/
    Update TS of x_i, by $TS = GTS$;
    Move x_i to the heads of LS_k and GS.
else
    /*A cache miss has occurred.*/
    if Cache is full **then**
        foreach block at the bottom of LS_j, $0 \leq j < n$ **do**
            if LS_j is a corresponding stack to a faulty disk **then**
                Its weight $W = GTS - TS$;
            else
                if It's in RAID-4 or RAID-5 **then**
                    Its weight $W = (GTS - TS) * (n - 1)$;
                else
                  /*Case of RAID-6*/
                  Its weight $W = (GTS - TS) * (n - 2)$;
        Delete the block with maximum W to obtain a free block;
    else
        /*Cache is not full.*/
        Get a free block.
    Load x_i to the free block.
    Update TS of x_i, by $TS = GTS$;
    Add x_i to the heads of GS and the corresponding LS.
Update GTS, by $GTS = GTS + 1$;
}

4.2 Detailed Description of VDF-LRU and VDF-LFU

Detailed descriptions of VDF-LRU and VDF-LFU for disk arrays of n-disk tolerating one or two disk failures are given in Algorithm 1 and Algorithm 2, respectively, using the variables summarized in Table 2.

5 SIMULATION RESULTS AND ANALYSIS

To evaluate the effectiveness of VDF, we conducted simulations under three typical workloads: SPC-1-web, LM-TBE, and DTRS.

SPC-1-web, a trace used in the SPC-1 benchmark suites, was collected in a search engine, which is widely used in the evaluation of storage systems [1], [11], [14]. LM-TBE and DTRS are provided by Microsoft Corporation collected in 2008. The LM-TBE trace was collected in back-end servers supporting a front-end Live Maps application. The DTRS trace was collected in a file server accessed by more than 3000 users to download various daily builds of Microsoft Visual Studio.

TABLE 2
Variants in VDF and Explanation

Variants	Explanation
x	A block request to buffer cache
LS	The local stack holding the blocks on one certain disk in the buffer cache
GS	The global stack holding all the blocks on all the disks in the buffer cache
n	The total number of disks including the faulty disk and surviving disks
TS	The timestamp of a block: records the access sequence number
F	The access frequency of a block
GTS	The global timestamp: it is equal to the timestamp of currently accessed block
W	The weight of a block

Algorithm 2: VDF-LFU for disk arrays of n disks tolerating one or two failures

```
Input: The request stream x₁, x₂, x₃, ..., xᵢ, ...
VDF_LFU_Replace(xᵢ){
/*For every i ≥ 1 and any xᵢ, one and only one of the following cases must
occur.*/
if xᵢ is in LSₖ, 0 ≤ k < n then
    /*A cache hit has occurred.*/
    Update F and TS of xᵢ, by F = F + 1;
    Move xᵢ to right place of LSₖ and GS according to F and TS.
else
    /*A cache miss has occurred.*/
    if Cache is full then
        foreach block at the bottom of LSⱼ, 0 ≤ j < n do
            if LSⱼ is a corresponding stack to a faulty disk then
                if It's in RAID-4 or RAID-5 then
                    Its weight W=F * (n − 1);
                else
                    /*Case of RAID-6*/
                    Its weight W=F * (n − 2);
            else
                Its weight W=F;
        Delete the block with minimum W and GTS − TS to obtain a free
        block;
    else
        /*Cache is not full.*/
        Get a free block.
    Load xᵢ to the free block.
    Initialize the frequency F and TS of xᵢ, by F = 1 and TS = GTS;
    Move xᵢ to right place of LSₖ and GS according to F and TS.
Update GTS, by GTS = GTS + 1;
}
```

Both traces were taken in a period of 24 hours and broken into pieces with 1-hour intervals [43]. We choose only the piece with most intensive I/O activities. For fairness and simplicity, we consider only the read operations and all block sizes are 4KB. Critical situation is only considered here that in our simulation we only focus on single disk failure in RAID-5 and concurrent double disk failures in RAID-6. We report RGR of LRU, LFU, VDF-LFU, and VDF-LRU under these workloads as shown in Figures 3, 4, 5, 6, 7, and 8, respectively.

Our simulator, named VDF-Sim, is written in C and the source code is approximately 3000 lines. It slices/splits the trace records into block requests as the input. Data blocks in a stack or blocks with the same hash values are linked via double circular lists. For a certain block in our simulator, we record its logical offset as the unique ID as the disk array is transparent to the upper level systems such as a file system. According to the data/parity distribution of the disk array and the logical offset of a block, it is easy to identify which disk the block resides on. The arriving timestamps of the requests are also recorded to evaluate p_i as we mentioned in Section

4, and to generate the misses trace in the next section.

The results show that, compared to the original LRU and LFU algorithms, VDF optimized algorithms achieved better performance consistently by reducing the RGR. Compared to LRU, VDF-LRU reduces the RGR by up to 31.4%, 36.2%, and 22.7% for RAID-5, and by up to 42.9%, 48.9%, and 28.3% for RAID-6 under SPC-1-web, LM-TBE, and DTRS traces, respectively. Compared to LFU, VDF-LFU reduces the RGR by up to 42.3%, 39.4%, and 24.4% for RAID-5, and by up to 50.2%, 50.7%, and 28.3% for RAID-6 respectively.

We find that the efficiency of VDF grows with the increased number of disks under the same number of cache-resident blocks in most cases. The efficiency of VDF is more significant with a moderate number of cache-resident blocks than that with a too small or too large number of cache-resident blocks. This can be explained as follows. The cache-resident blocks of a faulty disk in the original algorithm would occupy $1/n$ cache space with total n disks. With the fixed cache-resident blocks and the increased n, the number of cache-resident blocks of the faulty disk would be smaller. From the aspect of cache management, the impact of the marginal utility of blocks on hit ratio tends to decrease with the increased cache size. For example, adding $P1$ blocks to a cache with $P2$ blocks might improve the hit ratio with a larger gain compared to adding $P1$ blocks to cache with $P3$ blocks when $P2 < P3$. Thus, the marginal utility of blocks would be more obvious with more disks and thus the efficiency of VDF grows accordingly. However, if the number of cache-resident blocks is too small, it is hard to find hot blocks even with an extended period due to the large access interval. On the other hand, if the number of cache-resident blocks is too large, most of the requested blocks from the faulty disks would be cached, and the marginal utility of blocks becomes insufficient.

We also find that the VDF strategy becomes more efficient with LFU than LRU under the three workloads. One possible reason is that the temporal locality of these traces is weak as they are server-end traces and already filtered by upper level caches. Thus, the *Stack Depth Distribution* property of these traces is weak. As a result, the *Independent Reference Model* property of these traces would be relatively improved. Although we improve the weight of the blocks on faulty disk to $n − 1$ times for RAID-5 and $n − 2$ times for RAID-6 in both VDF-LRU and VDF-LFU, the efficiency is not the same. Mostly, the caching duration of blocks from a victim disk in VDF-LFU is longer than that in VDF-LRU, especially when the total number of cache-resident blocks is not large.

Simulation Results of the Web Trace for RAID-5

|	5 Disks						6 Disks						7 Disks						8 Disks					
	65536	131072	262144	524288	1048576	2097152	65536	131072	262144	524288	1048576	2097152	65536	131072	262144	524288	1048576	2097152	65536	131072	262144	524288	1048576	2097152
LRU	1.598	1.598	1.592	1.397	0.656	0.208	1.661	1.661	1.655	1.453	0.683	0.216	1.715	1.715	1.709	1.500	0.705	0.223	1.747	1.747	1.741	1.528	0.717	0.227
VDF-LRU	1.598	1.590	1.364	1.057	0.550	0.207	1.661	1.624	1.337	1.055	0.551	0.215	1.713	1.631	1.308	1.053	0.551	0.222	1.743	1.619	1.273	1.048	0.549	0.226
LFU	1.598	1.598	1.565	1.285	0.579	0.208	1.661	1.661	1.627	1.335	0.602	0.216	1.715	1.715	1.679	1.381	0.619	0.223	1.747	1.747	1.711	1.404	0.632	0.227
VDF-LFU	1.561	1.336	1.000	0.852	0.465	0.207	1.600	1.285	0.984	0.879	0.469	0.215	1.607	1.241	0.981	0.895	0.474	0.222	1.575	1.189	0.988	0.904	0.477	0.226

Fig. 3. Simulation results for a RAID-5 under the SPC-1-web trace. The number of disks ranges from 5 to 8, and the number of cache blocks varies from 64K or 65536 to 2M or 2097152 with the block size of 4KB.


Simulation Results of the LM-TBE Trace for RAID-5

	5 Disks						6 Disks						7 Disks						8 Disks					
	32768	65536	131072	262144	524288	1048576	32768	65536	131072	262144	524288	1048576	32768	65536	131072	262144	524288	1048576	32768	65536	131072	262144	524288	1048576
LRU	1.494	1.491	1.462	1.360	1.090	0.360	1.556	1.553	1.522	1.416	1.136	0.375	1.601	1.597	1.566	1.457	1.168	0.386	1.634	1.630	1.599	1.487	1.192	0.394
VDF-LRU	1.491	1.441	1.386	1.175	0.770	0.289	1.551	1.497	1.424	1.138	0.743	0.301	1.594	1.536	1.431	1.097	0.753	0.310	1.594	1.557	1.396	1.061	0.760	0.316
LFU	1.515	1.456	1.371	1.303	1.048	0.270	1.578	1.517	1.429	1.357	1.092	0.282	1.623	1.560	1.469	1.395	1.123	0.290	1.657	1.592	1.501	1.426	1.147	0.296
VDF-LFU	1.469	1.410	1.234	0.835	0.689	0.266	1.521	1.433	1.195	0.835	0.705	0.277	1.554	1.435	1.119	0.850	0.721	0.285	1.573	1.433	1.030	0.865	0.729	0.291

Fig. 4. Simulation results for a RAID-5 under the LM-TBE trace with various number of disks and cache blocks. The block size is 4KB.

Simulation Results of the DTRS Trace for RAID-5

|	5 Disks						6 Disks						7 Disks						8 Disks					
	32768	65536	131072	262144	524288	1048576	32768	65536	131072	262144	524288	1048576	32768	65536	131072	262144	524288	1048576	32768	65536	131072	262144	524288	1048576
LRU	1.562	1.530	1.475	1.311	0.824	0.639	1.628	1.594	1.537	1.367	0.860	0.666	1.672	1.637	1.579	1.405	0.883	0.684	1.706	1.672	1.613	1.434	0.901	0.699
VDF-LRU	1.545	1.484	1.353	1.083	0.766	0.639	1.603	1.527	1.325	1.097	0.787	0.666	1.638	1.549	1.293	1.106	0.800	0.684	1.665	1.560	1.291	1.109	0.814	0.699
LFU	1.547	1.489	1.406	1.211	0.900	0.639	1.613	1.552	1.465	1.263	0.939	0.666	1.656	1.594	1.505	1.298	0.964	0.684	1.690	1.627	1.538	1.326	0.986	0.699
VDF-LFU	1.480	1.367	1.149	0.970	0.792	0.639	1.522	1.372	1.153	1.000	0.820	0.666	1.523	1.370	1.154	1.023	0.838	0.684	1.537	1.359	1.163	1.041	0.853	0.699

Fig. 5. Simulation results for a RAID-5 under the DTRS trace with various number of disks and cache blocks. The block size is 4KB.

Simulation Results of the Web Trace for RAID-6

|	6 Disks						7 Disks						8 Disks						9 Disks					
	65536	131072	262144	524288	1048576	2097152	65536	131072	262144	524288	1048576	2097152	65536	131072	262144	524288	1048576	2097152	65536	131072	262144	524288	1048576	2097152
LRU	1.99799	1.99799	1.99044	1.74718	0.82109	0.26012	2.14103	2.14103	2.13277	1.87125	0.87844	0.27841	2.254	2.254	2.24542	1.97023	0.9246	0.29294	2.33034	2.33034	2.3215	2.03706	0.95673	0.30349
VDF-LRU	1.99799	1.98988	1.75843	1.1899	0.65673	0.25849	2.14103	2.12787	1.77119	1.182	0.66305	0.27655	2.254	2.23583	1.75457	1.17373	0.66707	0.29095	2.33025	2.26833	1.70815	1.16353	0.66826	0.30128
LFU	1.99799	1.9568	1.60771	0.72243	0.25964	2.14103	2.14103	2.09614	1.7201	0.77399	0.27794	2.254	2.254	2.20628	1.81132	0.81581	0.29246	2.33034	2.33034	2.2819	1.87224	0.8433	0.30293	
VDF-LFU	1.99799	1.90205	1.39714	0.90544	0.52449	0.25822	2.13251	1.95555	1.3451	0.90341	0.53845	0.27622	2.22249	1.95693	1.29612	0.9144	0.55044	0.29064	2.27686	1.91925	1.24024	0.9327	0.56049	0.30098

Fig. 6. Simulation results for a RAID-6 under the SPC-1-web trace. The number of disks ranges from 6 to 9, and the number of cache blocks varies from 64K or 65536 to 2M or 2097152 with the block size of 4KB.
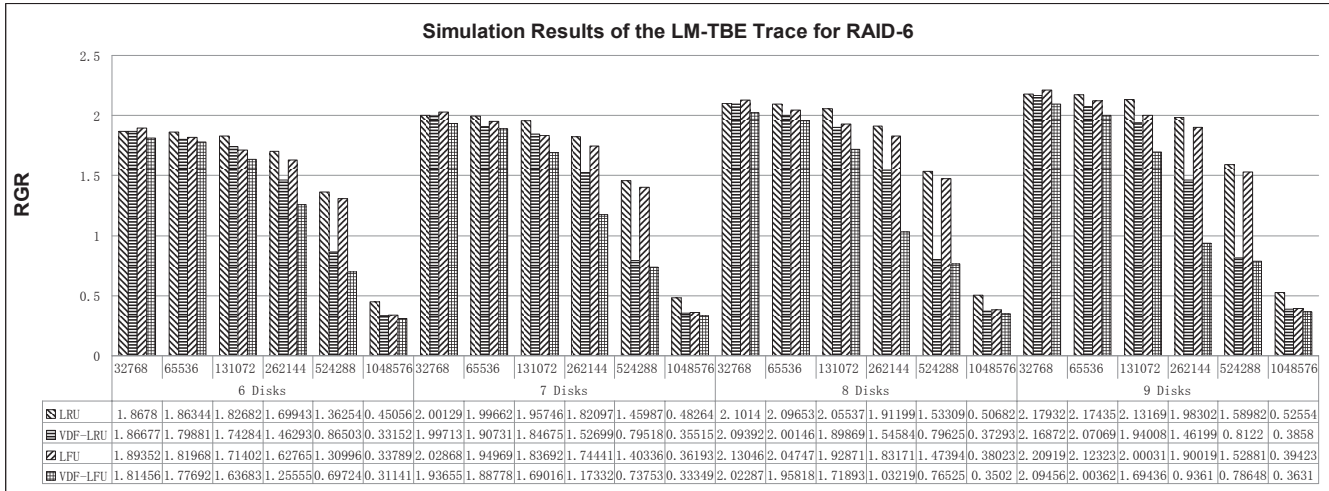

Simulation Results of the LM-TBE Trace for RAID-6

	6 Disks						7 Disks						8 Disks						9 Disks					
	32768	65536	131072	262144	524288	1048576	32768	65536	131072	262144	524288	1048576	32768	65536	131072	262144	524288	1048576	32768	65536	131072	262144	524288	1048576
LRU	1.8678	1.86344	1.82682	1.69943	1.36254	0.45056	2.00129	1.99662	1.95746	1.82097	1.45987	0.48264	2.1014	2.09653	2.05537	1.91199	1.53309	0.50682	2.17932	2.17435	2.13169	1.98302	1.58982	0.52554
VDF-LRU	1.86677	1.79881	1.74284	1.46293	0.86503	0.33152	2.09713	1.90731	1.84675	1.52699	0.79518	0.35515	2.09392	2.00146	1.89869	1.54584	0.79625	0.37293	2.16872	2.07069	1.94008	1.46199	0.8122	0.3858
LFU	1.89352	1.81968	1.71402	1.62765	1.30996	0.33789	2.02868	1.94969	1.83692	1.74441	1.40336	0.36193	2.13046	2.04747	1.92871	1.83171	1.47394	0.38023	2.20919	2.12323	2.00031	1.90019	1.52881	0.39423
VDF-LFU	1.81456	1.77692	1.63683	1.25555	0.69724	0.31141	1.93655	1.88778	1.69016	1.17332	0.73753	0.33349	2.02287	1.95818	1.71893	1.03219	0.76525	0.3502	2.09456	2.00362	1.69436	0.9361	0.78648	0.3631

Fig. 7. Simulation results for a RAID-6 under the LM-TBE trace with various number of disks and cache blocks. The block size is 4KB.

Simulation Results of the DTRS Trace for RAID-6

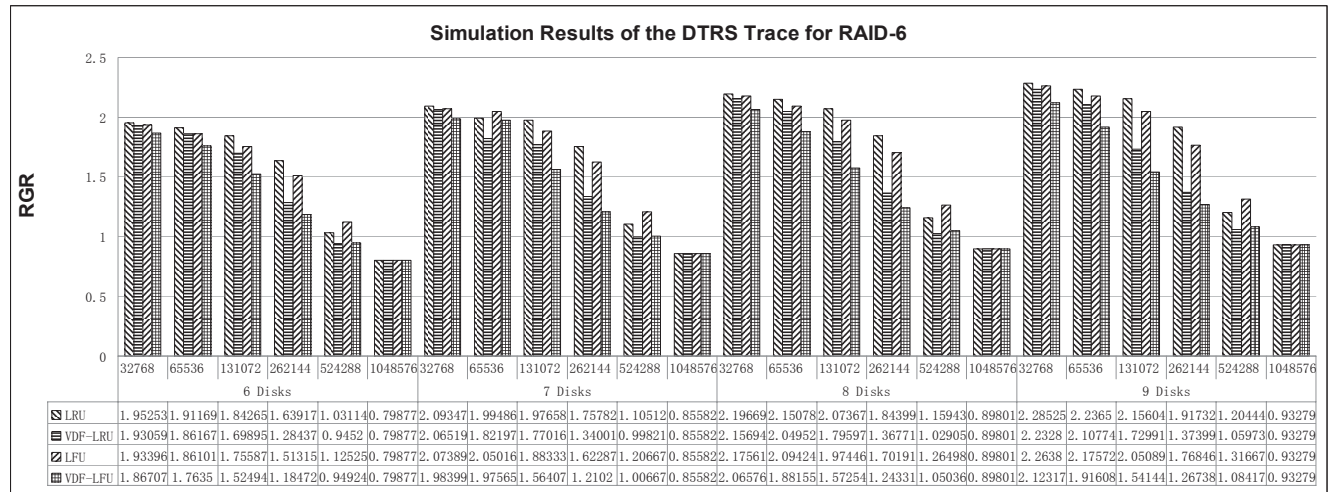|	6 Disks						7 Disks						8 Disks						9 Disks					
	32768	65536	131072	262144	524288	1048576	32768	65536	131072	262144	524288	1048576	32768	65536	131072	262144	524288	1048576	32768	65536	131072	262144	524288	1048576
LRU	1.95253	1.91169	1.84265	1.63917	1.03114	0.79877	2.09347	1.99486	1.97658	1.75782	1.10512	0.855822	2.19669	2.15078	2.07367	1.84399	1.15943	0.89801	2.28525	2.2365	2.15604	1.91732	1.20444	0.93279
VDF-LRU	1.93059	1.86167	1.69895	1.28437	0.9452	0.79877	2.06519	1.82197	1.77016	1.34001	0.99821	0.855822	2.15694	2.04952	1.79597	1.36771	1.029050	0.89801	2.2328	2.10774	1.72991	1.37399	1.059730	0.93279
LFU	1.93396	1.86101	1.75587	1.51315	1.12525	0.79877	2.073892	2.05016	1.88333	1.62287	1.20667	0.855822	2.17561	2.09424	1.97446	1.70191	1.264980	0.89801	2.2638	2.175722	2.05089	1.76846	1.31667	0.93279
VDF-LFU	1.86707	1.7635	1.52494	1.18472	0.949240	0.79877	1.98399	1.97565	1.56407	1.2102	1.00667	0.855822	2.065760	1.88155	1.572541	1.24331	1.050360	0.89801	2.12317	1.916081	1.54144	1.26738	1.08417	0.93279

Fig. 8.   Simulation results for a RAID-6 under the DTRS trace with various number of disks and cache blocks. The block size is 4KB.

Therefore, VDF-LFU works better than VDF-LRU in these traces in most cases.

From the simulation results, VDF cache seems to be more effective under the case of concurrent disk failures in RAID-6 than the case of single disk failure RAID-5. To our analysis, the reason is there are much more misses on faulty disks in cases of RAID-6 with concurrent double disk failures than that of RAID-5 with single disk failure. It would bring a higher RGR and leave a larger room for optimization.

6 PROTOTYPING OF VDF

To further evaluate VDF, we implemented a prototype of VDF in a software RAID system in Linux known as *MD*. In this section, we present our measurement results for both RAID-5 and RAID-6, including the efficiency of online recovery duration in full-bandwidth reconstruction mode and system service rate under the degraded mode without reconstruction.

6.1 Evaluation Methodology

Measurements on real world systems are welcome in research of computer systems. However, implementation in a real system is a lengthy process and always complex and challenging. Here, we use a straightforward and accurate measurement approach to evaluate the efficiency of VDF. The architecture of our prototype is described in Figure 9. First, we collect the cache miss information during our simulation in Section 5, which includes not only the block ID but also the real access timestamps. Then, we treat the RAID as a file device, and use an application in user mode to play the traces we have collected from our simulations which is similar to RAID-meter [1], [11]. However, the difference is that our application uses direct I/O (available in Linux -2.6 and up) instead of buffered I/O to avoid the requests being re-cached by the file system buffer cache. All missed I/O requests sent to the MD layer directly. Thus our simulation and the application join together to exploit the buffer cache and replacement algorithm. The trace player is also written in C and the source code is approximately 500 lines.

In our experiment, we evaluated the effectiveness of VDF, including the online reconstruction duration in full-bandwidth reconstruction mode, and the system serviceability in the degraded mode without reconstruction. For online reconstruction in full-bandwidth reconstruction mode, an open-loop measurement approach is adopted, where all filtered traces are played according to their timestamps as recorded in the original trace file. For degraded mode, a closed-loop measurement approach is adopted, where all filtered traces are played only according to their original sequence one by one and without any interval, to find the system serviceability.

6.2 Experimental Environment

In our experiment, we employ a SuperMicro storage server with two Intel(R) Xeon(R) X5560 @2.67GHz (6 cores) CPUs, 12GB DDR3 main memory. All disks are Western Digital WD10EALS Caviar Blue SATA2, which are connected by an Adaptec 31605 SAS/SATA RAID controller with a 256MB
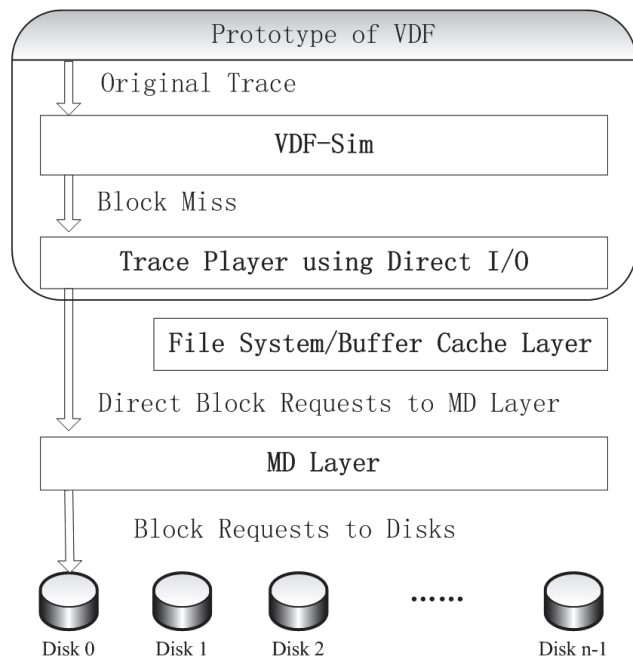
Fig. 9. Architecture of VDF prototype.

dedicated cache. We disabled the RAID function of the controller and only used the direct I/O mode to connect each disk. The operating system is Linux Fedora 12 X86_64 with the kernel version of 2.6.32.

In Linux, there is a software implementation of RAID called Multiple Devices *MD*, which is popular in verification of RAID optimization scheme [1], [11]. To facilitate the analysis and verification of VDF cache, we also used MD as our experimental platform. We used the default settings of MD: the chunk size is 64KB, the number of stripe-heads is 256 and the data layout is left-symmetric. In our open-loop testing, the minimum reconstruction bandwidth is set to 100MBps to utilize all remaining bandwidth for reconstruction besides the bandwidth taken by user workloads.

As VDF targets the storage server consisting of disk arrays which run under faulty conditions, we chose the server-end trace SPC-1-web as our experimental material, which spans a 60GB dataset. The workload is filtered by 131072 to 524288 blocks in our simulation to generate the experimental inputs. In the open-loop measurement, we test VDF with 5 to 8 disks. The results are reported in terms of reconstruction speed. The improvements of VDF over the original LRU and LFU are calculated. In the close-loop measurement, we used a multi-threaded application to play the filtered trace to measure the service rate. We also evaluated the impact of thread number to service rate, in addition to the impact of the number of blocks and the number of disks. The results are reported as system service rate improvement by VDF cache compared to original LFU and LRU. We ran each test three times and report the average. The results are very stable and consistent as the difference among all three rounds was very small (less than 5%).

TABLE 3
Experimental Results of an Open-loop Testing for RAID-5 of 5-8 Disks

Disks	Blocks	LRU (s)	VDF-LRU (s)	Improvement	LFU (s)	VDF-LFU (s)	Improvement
5 disks	131072	2662	2543	4.5%	2710	1929	28.8%
	262144	2958	1935	34.6%	2851	1531	46.3%
	524288	1845	1407	23.7%	1786	1310	26.7%
6 disks	131072	1176	1147	2.5%	1175	964	18.0%
	262144	1234	943	23.6%	1226	921	24.9%
	524288	1027	818	20.4%	1005	806	19.8%
7 disks	131072	730	685	6.2%	733	652	11.1%
	262144	758	659	13.1%	761	657	13.7%
	524288	691	599	13.3%	687	598	13.0%
8 disks	131072	504	485	3.8%	509	485	4.7%
	262144	558	501	10.2%	560	501	10.5%
	524288	527	483	8.4%	526	479	8.9%

TABLE 4
Experimental Results of an Open-loop Testing for RAID-6 of 6-9 disks.

Disks	Blocks	LRU (s)	VDF-LRU (s)	Improvement	LFU (s)	VDF-LFU (s)	Improvement
6 disks	131072	4667	4269	8.5%	4796	3843	19.8%
	262144	4447	2750	38.2%	4399	2323	47.2%
	524288	2567	1744	32.1%	2499	1623	35.1%
7 disks	131072	2180	2131	2.2%	2188	2024	4.9%
	262144	2168	1905	12.1%	2111	1676	20.6%
	524288	1773	1544	12.9%	1710	1470	14.0%
8 disks	131072	1510	1450	4.0%	1492	1340	10.2%
	262144	1504	1233	18.0%	1508	1106	26.7%
	524288	1305	1053	19.3%	1297	1038	20.0%
9 disks	131072	1103	1079	2.2%	1097	1056	3.7%
	262144	1102	1010	8.3%	1100	998	9.3%
	524288	1015	970	4.4%	1006	969	3.7%

6.3 Open-loop Measurement Results and Analysis

Table 3 and Table 4 describes the results under an open-loop testing using the SPC-1-web trace for RAID-5 and RAID-6 respectively, where the number of disks ranges from 5 to 8 in RAID-5 and from 6 to 9 in RAID-6, and the number of blocks ranges from 131072 to 524288. The experimental results of the open-loop testing show that compared to the original LRU and LFU algorithms, the VDF-optimized algorithm speeds up the online reconstruction process. The speedup of VDF-LFU is up to 46.3% in RAID-5 and 47.2% in RAID-6 compared to LFU. VDF-LRU speeds up the online reconstruction by up to 34.6% in RAID-5 and 38.2% in RAID-6 compared to LRU.

With the same number of cache-resident blocks, the experimental results show that the improvement of reconstruction durations of VDF-LFU to LFU decreases with the increased number of disks in most cases. Similar trend is also observed for the improvement of VDF-LRU over LRU, except for seven to eight disks in RAID-6 and low improvement of RGR using VDF when the number of blocks is 131072. These trends are in contrast to our previous simulation results where the efficiency of VDF tends to be more sufficient with increased number of disks. From Equation 4, the improvement of reconstruction durations of VDF-LFU to LFU, which is presented by $Imprv$, could be described with Equation 5. $BW_U \times RGR_{VDF}$ is always less than BW otherwise the system would be overloaded, so $\frac{BW}{RGR_{VDF}}$ is larger than BW_U. VDF works in most cases where $\frac{RGR_{ORI}}{RGR_{VDF}}$ is larger than 1, thus $\frac{BW_U * RGR_{ORI}}{RGR_{VDF}}$ is larger than BW_U. Therefore, the changing rate of improvement according to the number of total disks should only be determined by the changing rates of

$\frac{BW_U \times RGR_{ORI}}{RGR_{VDF}} - BW_U$ and $\frac{BW}{RGR_{VDF}} - BW_U$. Obviously, BW grows linearly with the total number of disks. The changing rate of the latter is larger than the change rate of BW. From the simulation result, we can find that the changing rate of the former is much smaller than that of BW with the number of disks from 5 to 8 in RAID-5. Therefore, in most cases shown above, the improvement of reconstruction durations of VDF cache to the original cache decreases with the increased number of disks. However, in case of RAID-6 with seven to eight disks, the changing rate increases to more than 20%. As a result, the improvement increases in such case.

$$RD_{Imprv} = \frac{\frac{BW_U * RGR_{ORI}}{RGR_{VDF}} - BW_U}{\frac{BW}{RGR_{VDF}} - BW_U} \tag{5}$$

With the same number of disks, we notice that the reconstruction duration of the trace filtered by 131072 blocks is less than the reconstruction duration of trace filtered by 262144 blocks in many cases. On one hand, as we use a number of blocks to warm up the cache, this part of the miss information is not recorded in our filtered trace file. The first 131072 block misses in the trace filtered by 262144 blocks has a lower average arrival rate than the remaining part. On the other hand, when the number of blocks in the cache is 131072 or 262144, the cache is too small to find the hot blocks, which implies fewer hits in those cases and the RGRs are similar. Therefore, the reconstruction duration of the trace filtered by 131072 blocks might be less than that of the trace filtered by 262144 blocks in those cases.

Fig. 10. Service rate improvement of VDF in degraded mode of a RAID-5 of 8 disks. The number of blocks is 524288 and the number of threads ranges from 20 to 80.


Fig. 11. Service rate improvement of VDF in degraded mode of a RAID-5 of 8 disks. The number of blocks is from 131072 to 524288 and the number of threads is 60.

Fig. 12. Service rate improvement of VDF in degraded mode of a RAID-5 of 5 to 8 disks. The number of blocks is 524288 and the number of threads is 60.

6.4 Close-loop Measurement Results and Analysis

Figures 10, 11, and 12 present the close-loop testing results under different scenarios with various numbers of threads, disks, and data blocks. The results are reported as service rate improvement, which is inversely proportional to the *Play Duration* (**PD**) of a whole filtered trace. For example, the corresponding service rate improvement of VDF-LRU to LRU should be calculated by $\frac{PD_{LRU}-PD_{VDF-LRU}}{PD_{VDF-LRU}}$.

From the experimental results, we find that VDF is effective in improving the system service rate. Compared to LFU, VDF-LFU improves the system service rate up to 46.8% with 60 threads under 8 disks and 262144 blocks. Compared to LRU, VDF-LRU improves the system service rate by up to 28.4% with 80 threads under 8 disks and 524288 blocks.

With the increasing number of I/O threads, the service rate improvement increases accordingly and gets close to the theoretical value calculated by the simulation results. Although the whole trace would be evenly distributed on all disks due to the round-robin addressing in RAID, the incoming user requests might not be evenly distributed on all the disks during a short period. Therefore, with a larger number of threads, which implies a longer scheduling window, the distribution of incoming user requests is more balanced and the user service rate is closer to the maximum system service rate.

With the same number of blocks and a fixed number of threads, the service rate improvement of VDF-LRU to LRU is consistent with the trend of the simulation result. However, to our surprise, the results of VDF-LFU to LFU were just opposite with the trend of the simulation result. As per our analysis, this was primarily due to two reasons. First, from the simulation result, the relative RGR reduction of VDF-LFU to LFU with 524288 blocks is in a small area from 33.7% to 35.6%. Second, the number of concurrent threads is fixed, which means that the number of threads per disk would increase with the decreased total number of disks. Thus, based on the above analysis, when the total number of disks is small, the improvement is closer to the theoretical value. Therefore, under close to theoretical peak service rates and more I/O threads per disk, the trend of service rate improvement of VDF-LFU to LFU is very possibly opposite with the trend of the simulation result. As a result, with the same number of disks and a fixed number of threads, which means a fixed number of I/O threads per disk, the service rate improvement is quite consistent with the trend of the simulation results.

6.5 Further Discussion

Several more issues deserve further discussion. The first issue is the implementation cost of VDF. As we mentioned in Section 4, to make the smooth conversion between the original cache algorithms and the VDF-based algorithms, two types of stacks should be employed to implement VDF cache. This adds both spatial and temporal overhead. The spatial overhead include the extra information in each block head such as the timestamp and the extra stack pointer of the local stack. Compared to the buffer cache size, this overhead is very small. The temporal overhead is the computation of the weight of block at the bottom of each LS stack. Due to the high computation ability of today's CPU, this should not influence the overall system performance.

Second, can we integrate VDF into other optimizations in faulty mode? As the VDF cache essentially reduces the user requests to the surviving disks, it can be integrated with other optimizations in faulty mode, such as optimization on data/parity/spare layout and reconstruction workloads. The

approach of redirection of reads utilizes the reconstructed data in a spare disk to serve part of the reads to the faulty disk. Thus the miss penalty of these reconstructed data block is zero in terms of extra requests to the surviving disks. There still exist hot data with large miss penalty on faulty disks. Therefore, VDF can still help.

Third, could RGR be suitable to describe the status of write operation? From its definition, RGR is determined by MP_i and p_i. The calculation of p_i in write operations is similar to read operations. However, the calculation of MP_i in write operations is quite different from read operations, as they might be done with two approaches based on the parity distribution in RAID with faulty disk(s). One is the Read Modify Write, and the other is Parity-Reconstruction-Write. Here, we take an example of short writes on an n-disk RAID-5 with one faulty disk to demonstrate the MP_i calculation for write operations. Once a short write is sent to the surviving disks and the corresponding parity is not on the faulty disk, the Read-Modify-Write should be performed which results in two reads and two writes on the surviving disks, and thus the MP_i is 4. Otherwise, the Parity-Reconstruction-Write should be performed which results in $n-1$ reads and one write on the surviving disks, and thus the MP_i is n.

7 CONCLUSIONS

In this paper, we present a penalty-aware buffer cache replacement strategy, named Victim (or faulty) Disk(s) First (VDF) cache, to improve the reliability and performance of a RAID-based storage system, particularly under faulty conditions. The basic idea of VDF is to treat the faulty disks more favorably, or give a higher priority to cache the data associated with the faulty disks. The benefit of this scheme is to reduce number of the cache miss directed to the faulty disk, and thus to reduce the I/O requests to the surviving disks overall. Less disk I/O activity caused by the user workflow will (1) improve the performance of the disk array, and (2) allow more bandwidth for online reconstruction which in turn speeds up the recovery, and thus improves the reliability. Our results based on both simulation and prototyping implementation has demonstrated the effectiveness of VDF in terms of reduced disk I/O activities and a faster recovery.

ACKNOWLEDGMENTS

REFERENCES

[1] S. Wu, H. Jiang, D. Feng, L. Tian, and Bo Mao. WorkOut: I/O workload outsourcing for boosting RAID reconstruction performance. In *Proceedings of the 7th USENIX FAST Conference*, USA, 2009.

[2] D. A. Patterson, G. A. Gibson, and R. H. Katz. A case for redundant arrays of inexpensive disks (RAID). In *Proceedings of the 1988 ACM SIGMOD international conference*, Chicago, Illinois, USA, June 1988.

[3] M. Holland and G. A. Gibson. Parity declustering for continuous operation in redundant disk arrays. In *Proceedings of the 5th Conference on Architectural Support for Programming Languages and Operating Systems*, pages 23–35, Boston, Massachusetts, USA, October 1992.

[4] J. Menon and D. Mattson. Distributed sparing in disk arrays. In *Proceedings of the 37th international conference on COMPCON*, pages 410–421, USA, 1992.

[5] Q. Xin, P. L. Miller, and T. J. E. Schwarz. Evaluation of distributed recovery in large-scale storage systems. In *Proceedings of 13th IEEE International Symposium on High Performance Distributed Computing*, pages 172–181, June 2004.

[6] G. K.M, X. Li, and J. J. Wylie. Flat XOR-based erasure codes in storage systems: Constructions, efficient recovery, and tradeoffs. In *26th IEEE MSST Conference*, USA, 2010.

[7] S. Wan, Q. Cao, C. S. Xie, B. Eckart, and X. He. Code-M: A non-MDS erasure code scheme to support fast recovery from up to two-disk failures in storage systems. In *IEEE/IFIP DSN International Conference*, USA, 2010.

[8] R. Y. Hou, J. Meno, and Y. N. Patt. Balancing I/O response time and disk rebuild time in a RAID5 disk array. In *Proceeding of the 26th Hawaii International Conference on System Sciences*, Jan 1993.

[9] J. Y.B. Lee and J. C.S. Lui. Automatic recovery from disk failure in continuous-media servers. *The Computer Journal*, 13(5), May 2002.

[10] M. Holland, G. A. Gibson, and D. P. Siewiorek. Fast, on-line failure recovery in redundant disk arrays. In *The 23rd International Symposium on Fault-Tolerant Computing*, pages 422–431, France, 1993.

[11] L. Tian, D. Feng, H. Jiang, K. Zhou, L. Zeng, J. Chen, Z. Wang, and Z. Song. PRO: A popularity-based multi-threaded reconstruction optimization for RAID-structured storage systems. In *Proceedings of the 5th USENIX FAST Conference*, pages 277–290, USA, 2007.

[12] L. Tian, H. Jiang, and D. Feng. Implementation and evaluation of a popularity-based reconstruction optimization algorithm in availability-oriented disk arrays. In *24th IEEE MSST Conference*, pages 233–238, USA, 2007.

[13] R. R. Muntz and J. C. S. Lui. Performance analysis of disk arrays under failure. In *Proceedings of the 16th International Conference on Very Large Databases*, pages 162–173, 1990.

[14] T. Xie and H. Wang. MICRO: A multilevel caching-based reconstruction optimization for mobile storage systems. *IEEE Transactions on Computers*, 57(10):1386–1398, Oct 2008.

[15] A. Dan and D. Towsley. An approximate analysis of the LRU and FIFO buffer replacement schemes. In *Proceedings of the 1990 ACM SIGMETRICS Conference*, pages 143–152, USA, 1990.

[16] J. T. Robinson and M. V. Devarakonda. Data cache management using frequency-based replacement. In *Proceedings ACM SIGMETRICS Conference*, pages 134–142, USA, 1990.

[17] E. J. O'Neil, P. E. O'Neil, and G. Weikum. The LRU-K page replacement algorithm for database disk buffering. In *Proceedings of the ACM SIGMOD international Conference*, pages 297–306, USA, 1993.

[18] E. J. O'Neil, P. E. O'Neil, and G. Weikum. An optimality proof of the LRU-K page replacement algorithm. *Journal of ACM*, 46(1):92–112, Jan 1999.

[19] T. Johnson and D. Shasha. 2Q: A low overhead high performance buffer management replacement algorithm. In *Proceedings of the 20th International Conference on Very Large Databases*, pages 439–450, USA, 1994.

[20] D. Lee, J. Choi, J. H. Kim, S. H. Noh, S. L. Min, Y. Cho, and C. S. Kim. On the existence of a spectrum of policies that subsumes the least recently used (LRU) and least frequently used (LFU) policies. In *Proceedings of the ACM SIGMETRICS Conference*, USA, 1994.

[21] D. Lee, J. Choi, J. H. Kim, S. H. Noh, S. L. Min, Y. Cho, and C. S. Kim. LRFU: A spectrum of policies that subsumes the least recently used and least frequently used policies. *IEEE Transactions on Computers*, 50(12):1352–1361, 2001.

[22] Y. Zhou, J. F. Philbin, and K. Li. The multi-queue replacement algorithm for second level buffer caches. In *Proceedings of 2001 USENIX Annual Technical Conference*, USA, 2001.

[23] Y. Zhou, J. F. Philbin, and K. Li. Second-level buffer cache management. *IEEE Transactions on Parallel and Distributed Systems*, 15(6):505–519, June 2004.

[24] S. Jiang and X. Zhang. LIRS: An efficient low inter-reference recency set replacement policy to improve buffer cache performance. In *Proceedings of the 2002 ACM SIGMETRICS Conference*, pages 31–42, USA, 2002.

[25] Song Jiang and Xiaodong Zhang. Making LRU friendly to weak locality workloads: a novel replacement algorithm to improve buffer cache performance. *IEEE Transactions on Computers*, 54(8):939–952, Aug 2005.

[26] N. Megiddo and D. S. Modha. ARC: A self-tuning, low overhead replacement cache. In *Proceedings of the 2nd USENIX FAST Conference*, pages 115–130, USA, 2003.

[27] S. Jiang, X. Ding, F. Chen, E. Tan, and X. Zhang. DULO: An effective buffer cache management scheme to exploit both temporal and spatial localities. In *Proceedings of the 4th USENIX FAST Conference*, San Francisco, CA, USA, 2005.

[28] X. Ding, S. Jiang, F. Chen, K. Davis, and X. Zhang. DiskSeen: Exploiting disk layout and access history to enhance I/O prefetch. In *Proceedings of the 2007 USENIX Annual Technical Conference*, USA, 2007.

[29] SNIA. Raid6 definition in dictionary r. http://www.snia.org/education/dictionary/r, 2007.

[30] Irving S Reed and G. Solomon. Polynomial codes over certain finite fields. *Journal of the Society for Industrial and Applied Mathematics*, 8:300–304, June 1960.

[31] Johannes Blomer, Malik Kalfane, Richard Karp, Marek Karpinski, Michael Luby, and David Zuckerman. An XOR-based erasure-resilient coding scheme. Technical Report TR-95-048, International Computer Science Institute, Berkeley, California, August 1995.

[32] Mario Blaum, Jim Brady, Jehoshua Bruck, and Jai Menon. EVENODD: An efficient scheme for tolerating double disk failures in raid architectures. *IEEE Transactions On Computers*, 44(2):192–202, February 1995.

[33] Peter Corbett, Bob English, Atul Goel, Tomislav Grcanac, Steven Kleiman, James Leong, and Sunitha Sankar. Row-diagonal parity for double disk failure correction. In *Proceedings of the 3rd USENIX FAST Conference*, pages 1–14, USA, March 2004.

[34] Mario Blaum and Ron M. Roth. On lowest density MDS codes. *IEEE Transactions on Information Theory*, 45(1):46–59, January 1999.

[35] James S. Plank. The RAID-6 liberation codes. In *Proceedings of the 6th USENIX FAST Conference*, pages 97–110, San Jose, February 2008.

[36] Yuval Cassuto and Jehoshua Bruck. Cyclic lowest density MDS array codes. *IEEE Transactions on Information Theory*, 55(4):1721–1729, April 2009.

[37] Lihao Xu and Jehoshua Bruck. X-Code: MDS array codes with optimal encoding. *IEEE Transactions on Information Theory*, 45(1):272–276, January 1999.

[38] Chao Jin, Hong Jiang, Dan Feng, and Lei Tian. P-code: A new RAID-6 code with optimal properties. In *ICS-2009: Proceedings of the 23rd International Conference on Supercomputing*, pages 360–369, Yorktown Heights, NY, USA, June 2009.

[39] ChanIk Park. Efficient placement of parity and data to tolerate two disk failures in disk array systems. *IEEE Transactions on Parallel and Distributed Systems*, 6(11):1177–1184, November 1995.

[40] James Lee Hafner. HoVer erasure codes for disk arrays. In *In DSN-2006: the 36th Annual IEEE/IFIP DSN International Conference*, pages 217–226, PA, June 2006.

[41] C. Wu, S. Wan, X. He, Q. Cao, and C. Xie. H-Code: A hybrid MDS array code to optimize partial stripe writes in raid-6. In *25th IEEE International Parallel and Distributed Processing Symposium*, USA, 2011.

[42] C. Wu, X. He, G. Wu, S. Wan, X. Liu, Q. Cao, and C. Xie. HDP code: A horizontal-diagonal parity code to optimize I/O load balancing in raid-6. In *Proceeding of 41th IEEE DSN International*, China, June 2011.

[43] S. Kavalanekar, B. Worthington, Q. Zhang, and V. Sharda. Characterization of storage workload traces from production windows servers. In *IEEE International Symposium on Workload Characterization*, USA, 2008.

[44] Shenggang Wan, Qiang Cao, Jianzhong Huang, Shiyi Li, Xin Li, Shenhui Zhan, Changsheng Xie, Li Yu, and Xubin He. Victim disk first: An asymmetric cache to boost the performance of disk arrays under faulty conditions. In *Proceedings of the USENIX Annual Technical Conference 2011*, Portland, OR, USA, June 2011.

Shenggang Wan received his PhD and BS degrees in computer science and technology in 2010 and 2003, MS degree in software engineering in 2005, all from Huazhong University of Science and Technology (HUST), China. He is currently a Post Doctoral researcher in the Department of Electronics and Information Engineering at HUST and a visiting scholar at Virginia Commonwealth University. His research interests include dependable storage systems and coding theory.

Xubin He received the PhD degree in electrical engineering from University of Rhode Island, USA, in 2002 and both the BS and MS degrees in computer science from Huazhong University of Science and Technology, China, in 1995 and 1997, respectively. He is currently an associate professor in the Department of Electrical and Computer Engineering at Virginia Commonwealth University. His research interests include computer architecture, storage systems, virtualization, and high availability computing. He received the Ralph E. Powe Junior Faculty Enhancement Award in 2004 and the Sigma Xi Research Award (TTU Chapter) in 2005 and 2010. He is a senior member of the IEEE, a member of the IEEE Computer Society, ACM, and Sigma Xi.

Jianzhong Huang received the PhD degree in computer architecture in 2005 and completed the Post Doctoral research in information engineering in 2007 from Huazhong University of Science and Technology(HUST). He is currently an associate professor in the Department of Computer Engineering at HUST. His research interests include computer architecture and networked storage systems. He is a member of China Computer Federation (CCF).

Qiang Cao received the PhD degree in computer architecture and MS degree in computer technology in 2003 and 2000, and BS in applied physics in 1997. He is currently a professor at Huazhong University of Science and Technology. His research interests include computer architecture, large scale storage systems, and performance evaluation on computer systems. He is a senior member of China Computer Federation (CCF) and a member of the IEEE.

Shiyi Li received the BS in mathematics in 2009. He is currently a PHD student at Huazhong University of Science and Technology. His research interests include dependable storage systems and coding theory.

Changsheng Xie received the BS and MS degrees in computer science both from Huazhong University of Science and Technology, in 1982 and 1988, respectively. He is currently a professor in the Department of Computer Engineering at HUST. He is the deputy director of the Wuhan National Laboratory for Optoelectronics. His research interests include computer architecture, disk I/O system, networked data storage system and digital media technology. He is the vice chair of the expert committee of Storage Networking Industry Association (SNIA), China. He is a member of the IEEE.