

ETICA: Efficient Two-Level I/O Caching Architecture for Virtualized Platforms

Saba Ahmadian, Reza Salkhordeh, Onur Mutlu, Hossein Asadi

Abstract—In recent years, increased I/O demand of *Virtual Machines* (VMs) in large-scale data centers and cloud computing has encouraged system architects to design high-performance storage systems. One common approach to improving performance is to employ fast storage devices such as *Solid-State Drives* (SSDs) as an I/O caching layer for slower storage devices. SSDs provide high performance, especially on random requests, but they also have limited endurance: they support *only* a limited number of write operations and can therefore *wear out* relatively fast due to write operations. In addition to the write requests generated by the applications, each read miss in the SSD cache is served at the cost of imposing a write operation to the SSD (to copy the data block into the cache), resulting in an even larger number of writes into the SSD. Previous I/O caching schemes on virtualized platforms *only* partially mitigate the endurance limitations of SSD-based I/O caches; they mainly focus on assigning efficient cache write policies and cache space to the VMs. Moreover, existing cache space allocation schemes have inefficiencies: they *do not* take into account the impact of cache write policy in reuse distance calculation of the running workloads and hence, reserve cache blocks for accesses that would *not* be served by cache.

In this paper, we propose an *Efficient Two-Level I/O Caching Architecture* (ETICA) for virtualized platforms that can significantly improve I/O latency, endurance, and cost (in terms of cache size) while preserving the reliability of write-pending data blocks. As opposed to previous *one-level* I/O caching schemes in virtualized platforms, our proposed architecture 1) provides *two* levels of cache by employing both *Dynamic Random-Access Memory* (DRAM) and SSD in the I/O caching layer of virtualized platforms and 2) effectively partitions the cache space between running VMs to achieve maximum performance and minimum cache size. To manage the two-level cache, unlike the previous reuse distance calculation schemes such as *Useful Reuse Distance* (URD), which only consider the request type and neglect the impact of *cache write policy*, we propose a new metric, *Policy Optimized reuse Distance* (POD). The key idea of POD is to effectively calculate the reuse distance and estimate the amount of two-level DRAM+SSD cache space to allocate by considering both 1) the request type and 2) the cache write policy. Doing so results in enhanced performance and reduced cache size due to the allocation of cache blocks *only* for the requests that would be served by the I/O cache. ETICA maintains the reliability of write-pending data blocks and improves performance by 1) assigning an effective and fixed write policy at each level of the I/O cache hierarchy and 2) employing effective promotion and eviction methods between cache levels. Our extensive experiments conducted with a real implementation of the proposed two-level storage caching architecture show that ETICA provides 45% higher performance, compared to the state-of-the-art caching schemes in virtualized platforms, while improving both cache size and SSD endurance by 51.7% and 33.8%, respectively.

1 INTRODUCTION

Virtualized platforms are widely used in large scale data centers to provide significantly improved availability and flexibility. In such platforms, multiple *Virtual Machines* (VMs) provide different services on shared hardware. VM management and resource partitioning between VMs are performed by the hypervisor, which plays a major role in the virtualized platforms and aims to maximize performance and system utilization. The key advantages of virtualized platforms that are continuously evolving in computing industry are: 1) high flexibility due to the ability to run multiple VMs with different *Operating Systems* (OS), 2) high resource utilization, 3) resource isolation between different VMs, and 4) allocation of dynamically adjustable resources to the VMs [1].

The growing popularity of I/O intensive applications in data centers, such as *Online Transaction Processing* (OLTP), banking, data analysis, and other big data workloads, greatly increases the demand for high-performance storage subsystems. Existing storage systems still mainly consist of high-capacity and low-performance *Hard Disk Drives* (HDDs) with the goal of storing very large amounts of data at low cost. Such storage systems have become the performance bottleneck in large-scale data centers due to the growing performance gap between HDDs and processing elements in enterprise servers (as depicted by *I/O Operations Per Second* (IOPS) performance in Fig. 1).

To alleviate the performance shortcomings of such HDD-based storage systems, enterprise manufacturers, such as Dell EMC, HP, and NetApp [3], [4], [5], and various research works on emerging storage architectures, such as [6], [7], [8], [9], [10], [11], [12], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22], [23], propose *I/O caching* schemes that accelerate the response

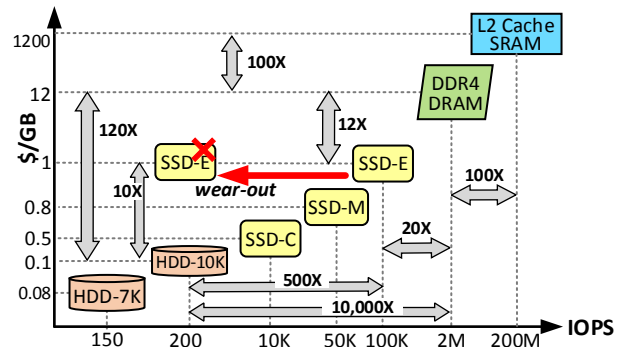


Fig. 1: Comparison of different storage devices in terms of cost (\$/GB) and performance (IOPS). SSD-C: Consumer SSD, SSD-M: Midrange SSD, SSD-E: Enterprise SSD, DDR4 DRAM: Double Data Rate 4th generation, HDD-7K: HDD with 7K *Revolutions Per Minute* (RPM), HDD-10K: HDD with 10K RPM [2].

time of I/O requests. High-performance storage devices such as *Dynamic Random Access Memories* (DRAMs) and *Solid-State Drives* (SSDs) can be employed in the I/O caching layer to alleviate the low performance of HDDs. As depicted in Fig. 1, DRAM provides the highest performance among many types of storage/memory devices [24] and thus seems to be the best candidate to employ in the I/O caching layer. However, DRAM is expensive, and using *volatile* DRAM as a cache comes with reliability hurdles for write requests. A DRAM-based I/O cache requires additional battery backups to correctly maintain the data buffered in DRAM to survive

power outages and system failures. Compared to DRAM, enterprise SSDs have about 12X lower cost and offer 20X less performance but they provide 500X higher performance than HDDs. In contrast to DRAM, SSDs are non-volatile and typically do not suffer from data loss due to power outages, but they support *only* a limited number of reliable writes [25], [26], [27], [28], [29], [30]. Thus, the main shortcoming of SSDs is their endurance limitation where both performance and reliability are significantly degraded when the number of committed writes exceeds the endurance limit (i.e., when the SSD wears out, as shown in Fig. 1). In addition, frequent replacement of such expensive SSDs imposes significant cost and reliability issues¹ in storage systems. Recent studies, such as [6], [7], [8], [9], [10], [11], [12], [13], [14], [15], [16], [17], [18], [19], [20], [32], propose I/O caching schemes based on SSDs. Among such studies, only few are applicable to virtualized platforms [6], [9], [10], [11], [32]. Such I/O caching schemes aim to either only improve performance without considering the endurance of the SSD [7], [10], [11], [12], [13], [14], [32] or attempt to improve performance with minimum overhead on SSD endurance [6], [8], [9]. In addition, such studies suffer from two main shortcomings: 1) their performance improvement is limited by the high SSD latency (relative to DRAM) and 2) their endurance-preserving techniques lead to performance degradation.

Another approach to improving the performance of storage systems is to use hybrid techniques that take advantage of *both* DRAM and SSDs in the I/O caching layer. Multiple industrial approaches (e.g., the Dell EMC FAST cache [3] and ZFS L2ARC cache [33], [34], [35], [36]) and academic studies [37], [38], [39], [40], [41], [42], [43], [44], [45], [46], [47], [48], [49], [50], [51], [52], [53] employ two levels of I/O caching using both DRAM and emerging *Non-Volatile Memories* (NVMs)/SSDs in their storage architectures. Such caching schemes, however, suffer from several important shortcomings: 1) They cannot provide cache space partitioning across VMs, which results in poor performance and high cost in virtualized platforms. 2) Buffered data in DRAM cache is at risk of loss due to sudden power outage; additional battery backups add significant cost overhead to the caching scheme. 3) The employed mechanisms for detecting hot data blocks are not accurate since they do not consider any workload characteristics, such as reuse distance of the accesses. 4) Past works do not provide write policy management at DRAM or SSD levels and most schemes configure both cache levels to use the *Write Back* (WB)² policy. 5) Such schemes only focus on detecting and promoting *hot* data blocks to the cache while there is no specified method to detect and evict *cold* data blocks (i.e., least recently accessed data blocks) from cache.

Most recent cache space partitioning schemes in virtualized platforms employ reuse distance analysis to estimate each VM's cache size, but they neglect key parameters such as *request type* and *cache write policy* [6], [11]. The state-of-the-art scheme, ECI-Cache [6], proposes *Useful Reuse Distance* (URD) and considers the *request type* in reuse distance calculation. However, it neglects the impact of *cache write policy* on the reuse distance calculation, and, as a result, it over-estimates the cache sizes for caches with write policies other than WB and *Write Through* (WT). ECI-Cache estimates the cache sizes of the VMs based on the URD metric and assigns a cache write policy on a per-VM basis. As such, it results in improved SSD endurance and performance-per-cost. The URD metric

is optimized for caches with WB and WT policies. It *over-estimates* the cache size for caches with other write policies (e.g., *Read Only* (RO)). If we use the URD metric in a cache with the RO policy, cache blocks would be allocated for *write* requests that would not be buffered in the RO cache, and thus the allocated blocks remain *unused*.

In this paper, we first propose a new metric, called *Policy Optimized reuse Distance* (POD), which considers the impact of the *cache write policy* in addition to *request type* in reuse distance calculation. POD, unlike URD, does *not* reserve cache blocks for accesses that would not be supplied by the cache and, hence, allocates much smaller cache space for a VM, compared to URD, while preserving the VM's performance. Second, we propose the *Efficient Two-Level I/O Caching Architecture* (ETICA) for virtualized platforms. ETICA employs *two levels* of cache, with DRAM at the first and SSD at the second level. Using the POD metric, the proposed architecture effectively partitions the space of both levels of cache between VMs and improves both performance and endurance.

In our proposal, we first assign an effective and fixed write policy to each cache level that provides high reliability and high endurance in DRAM and SSD levels, respectively. Second, using the POD metric, we partition the cache space by assigning an efficient cache size to each VM, which maximizes performance-per-cost. Third, we propose promotion and eviction methods to effectively transfer *popular* (i.e., *hot*) and *unpopular* (i.e., *cold*) data blocks between the two cache levels. Unlike previous I/O caching schemes, our proposed two-level cache does *not* evict popular data blocks from the cache until they become unpopular. Via online monitoring of I/O characteristics of the running workload, we extract several metrics such as 1) frequency, 2) type, and 3) POD of accesses in the running VMs. The extracted information is used to 1) estimate an efficient cache size for each VM and 2) detect *popular* and *unpopular* data blocks. To summarize, our proposed two-level I/O caching scheme aims to: 1) enhance both read and write performance, 2) overcome the reliability issues of the volatile DRAM cache, 3) improve the lifetime of the SSD cache, and 4) reduce the overall cost of the two-level cache (in terms of cache size allocated to each VM).

We implement our proposed two-level I/O caching mechanism on a real virtualized system including: 1) per-VM cache size estimation, 2) write policy assignment, 3) popular/unpopular data block detection, and 4) promotion/eviction mechanisms. We evaluate ETICA with more than 10 VMs running real application workloads from SNIA [54], [55]. Our experimental results show that the proposed architecture provides 45% lower I/O latency while also improving effective cache size and SSD endurance by 51.7% and 33.8%, respectively, compared to the best state-of-the-art caching scheme in virtualized platforms.

We make the following contributions.

- We propose a novel *two-level I/O cache* for virtualized platforms, using DRAM and SSD, that resolves reliability, endurance, and performance issues *without* any additional high-cost peripherals such as battery backups or internal disks.
- We propose a new metric, *Policy Optimized reuse Distance* (POD), which refines the concept of useful reuse distance calculation based on *cache write policy*. This metric does *not* allocate cache space for accesses that would not be served by the cache. Thus, it assigns a smaller cache space to each VM, resulting in lower cost.
- We propose a mechanism to assign efficient write policies for different cache levels to balance endurance and reliability of the I/O cache. Our mechanism enables read

1. Due to the degraded mode of *Redundant Array of Independent Disks* (RAID) [31] during the rebuild process in SSDs.

2. In a WB cache, the data block is written into the cache but the corresponding location in the hard disk is updated only when the data block is evicted from the cache, or at specific time intervals.

requests to be supplied from DRAM and write requests to be supplied from the SSD.

- Via online characterization of I/O requests, our technique effectively determines *popular* and *unpopular* data blocks. Our proposed promotion and eviction methods buffer and hold data blocks in the cache while they remain *popular*.
- We implement the proposed two-level I/O cache in a real system and evaluate the performance and endurance of our scheme by performing extensive experiments with more than 10 VMs running real application workloads on an open-source hypervisor, QEMU [56].
- We find ETICA to provide higher performance, higher efficiency, and higher endurance than the state-of-the-art I/O caching schemes for virtualized platforms [6].

2 RELATED WORK

In this section, we first describe and analyze the previous I/O caching architectures in virtualized platforms. Second, we analyze the existing I/O cache architectures in two groups: 1) single level and 2) multi level, which are mainly employed in *non-virtualized* platforms.

2.1 I/O Caching in Virtualized Platforms

S-CAVE [10], vCacheShare [9], Centaur [11], and ECI-Cache [6] are the most state-of-the-art hypervisor-based I/O caching schemes in virtualized platforms. Such schemes employ a *single level* of SSD cache and mainly focus on dynamic and efficient cache space partitioning between running VMs. S-CAVE [10] uses the *Working Set Size* (WSS) of the VMs for cache space estimation. To preserve the reliability of write requests, the write policy of the SSD cache is set to *Write Through* (WT). Such cache size estimation (based on WSS) is deprecated and fails to accurately estimate the cache space needed for workloads with sequential access patterns [6].

vCacheShare [9] estimates the cache size of the VMs based on locality and reuse intensity (i.e., workloads' burstiness). vCacheShare considers both reliability and endurance by applying *Write Around* (WA) or *Read Only* (RO) write policies which direct the write requests to the disk subsystem and *only* improve the performance of read accesses via caching. The cache size estimation scheme used in vCacheShare is applicable to CPU caches, but its assumptions (e.g., reuse intensity) *cannot* be applied in I/O caches [11]. vCacheShare does *not* improve the performance of write requests.

Centaur [11] assigns cache space to the running VMs based on reuse distance analysis, which is an effective approach to cache size estimation. The write policy of the SSD cache is simply set to *Write Back* (WB). Such a cache size estimation scheme does not consider either the *request type* or the *cache write policy*, which leads to over-estimation of the cache sizes, negatively impacting cost. In addition, by using the WB policy, Centaur negatively affects both storage reliability and the endurance of the SSD cache.

ECI-Cache [6] is the latest state-of-the-art I/O caching scheme for virtualized platforms, which overcomes the shortcomings of previous schemes. It proposes the *Useful Reuse Distance* (URD) metric, which considers *request type* in cache size estimation and reduces the allocated cache space while preserving performance, enabling lower cost. ECI-Cache dynamically assigns either the WB or the RO policies for each VM and thus provides higher performance and endurance for the SSD cache. There are four issues with ECI-Cache, which we aim to overcome with our new design. **First**, ECI-Cache dynamically assigns RO and WB policies on the VM's cache, while URD-based cache size estimation does *not* consider the *cache write policy* and only works for WB and WT caches.

URD reserves cache blocks for the accesses that would *not* be served by the cache due to different cache write policies, and therefore, it *over-estimates* the size of the caches using other policies, such as RO. For instance, in an RO cache, URD reserves cache blocks for write accesses that would *not* be buffered in the cache. **Second**, ECI-Cache is only a *one-level* (i.e., SSD-only) I/O cache where the performance improvement is limited by the performance of the SSD. **Third**, ECI-Cache updates the cache content on each cache miss and therefore imposes a large number of unnecessary writes into the SSD cache. **Fourth**, ECI-Cache's cache update scheme is supposed to promote any missed data blocks into the cache. Such scheme may evict a *hot* data block to promote a data block without any future references into the cache, resulting in performance degradation.

To summarize, existing hypervisor-based I/O caching architectures suffer from three major shortcomings: 1) they employ only an SSD in the caching layer, and thus their performance improvement is limited by the performance of the SSD (which is much lower than that of DRAM), 2) they fail in cache size estimation under different workloads and different cache write policies, and 3) they fail in balancing performance and endurance.

2.2 I/O Caching in Non-Virtualized Platforms

2.2.1 Single Level I/O Caches

A relatively high-performance memory device, such as DRAM or SSD is employed in the caching layer to close the performance gap between the processing units and the storage subsystem (which is mainly composed of low-performance HDDs) [7], [8], [12], [14], [18], [19], [57]. Such I/O caching schemes improve the performance of I/O requests by providing efficient cache space allocation and configuration. However, they neglect other key parameters such as storage reliability and endurance of the SSDs. Argon [57] aims to efficiently partition the DRAM cache space between different services to maximize the cache hit ratio. Janus [12] allocates SSD space based on the ratio of read accesses of the workload and aims to improve the performance of read requests. Hystor [7] employs an SSD cache to improve the performance of read and write accesses and aims to identify and buffer the data blocks, which helps in improving the hit ratio. ReCA [8] characterizes the workloads of running applications and provides a per-application cache configuration. This scheme aims to improve both performance and endurance by assigning different cache policies to different applications, but it cannot minimize the number of SSD writes (mainly the writes due to read misses) because it uses a single SSD in the caching layer. ReCA is able to reconfigure 1) cache line size, 2) write policy, and 3) eviction policy. This scheme allocates cache space globally and neglects the cache management in case of running multiple workloads (i.e., services). Such scheme cannot be applied in virtualized platforms, where multiple VMs are running on the system. Since this approach only employs SSDs in the caching layer, the performance improvement is limited by the highest performance of the SSDs. SHARDS [58] presents a hashed approximate reuse distance sampling scheme to be used in cache size estimation. This scheme improves the performance of reuse distance calculation without considering requests type and cache write policy in reuse distance calculation.

2.2.2 Multi Level I/O Caches

Such schemes employ both SSD and DRAM in the I/O caching layer.³ uCache [37] provides a simple two-level I/O cache design by employing DRAM and SSD. In this scheme, data blocks are kept in DRAM until they are evicted, and after eviction they are moved to the SSD. uCache improves the process of demoting data blocks from DRAM to SSD by aggregating a large number of small writes and, hence, improves the endurance of the SSD. However, write-pending data blocks are kept in the *volatile* DRAM cache, which results in data loss in case of power failure. Molar [38] presents a simulator-based hybrid I/O cache without any improvement in eviction/promotion methods. This scheme consists of DRAM, SSD, and HDD tiers and migrates data blocks between them. Molar aims to predict the future accesses and decides which data block should be evicted from DRAM based on this prediction. This scheme does *not* employ any cache write policy management on different tiers and hence cannot preserve 1) the reliability of buffered data blocks in the volatile DRAM cache, and 2) the endurance of the SSD cache.

In [22], DRAM, *Write Optimized SSD* (WO-SSD), and *Read Optimized SSD* (RO-SSD) are used in the caching layer. Write requests are buffered in both DRAM and WO-SSD, read misses are promoted to DRAM, and evictions from DRAM are directed into the RO-SSD. The main shortcomings of this architecture are: 1) Write requests are directed to both DRAM and WO-SSD, hence, they experience the write latency of the SSD. Using DRAM to buffer the write accesses has *no* impact on the performance of writes and only helps to improve the latency of future reads that may hit in DRAM (i.e., RAW accesses). The limited DRAM space and small ratio of RAW accesses lead to very small performance improvement by this architecture. 2) This architecture prevents buffering write accesses in the RO-SSD (to preserve the lifetime of RO-SSD), but on the other hand, evictions from the DRAM and promotions of read misses are performed on RO-SSD which impose extra writes and have a negative impact on the limited lifetime of RO-SSDs. 3) This method is a *global* cache which cannot assign efficient cache space for the running services, and hence, is not applicable in virtualized platforms.

Zettabyte File System (ZFS) *Level 2 Adaptive Replacement Cache* (L2ARC) [33], [34], [35], [36] is a file system level cache, which improves upon *Adaptive Replacement Cache* (ARC) [69] by employing an SSD between DRAM and HDD (i.e., disk subsystem), thereby reducing the latency of read misses. L2ARC works in a simple *First In First Out* (FIFO) mode and only improves the performance of read requests. As shown in Fig. 2a⁴, read requests may be supplied from DRAM or SSD (i.e., hit), and read misses are supplied from the disk subsystem and promoted into DRAM. L2ARC predicts to-be-evicted data blocks from DRAM and pushes them to the SSD (before they have to be evicted). The major shortcomings of L2ARC are: 1) it lacks an efficient promotion/eviction method: only a simple FIFO manages the contents of the SSD, 2) it could waste the SSD space by pushing data blocks from DRAM to SSD before their eviction, 3) by keeping evicted

3. Many recent works examine DRAM cache design [59], [60], [61], [62], [63], [64], [65] for hybrid main memory systems, e.g., those combining DRAM and *Phase Change Memory* (PCM) [59], [66], [67] or *Spin-Transfer Torque Magnetic Random-Access Memory* (STT-MRAM) [68]. Our work is significantly different, since none of these recent works 1) are applicable to I/O caching in virtualized storage systems (as they focus on main memory), 2) develop the new reuse distance metric we introduce and evaluate. As such, our work is orthogonal to such hybrid main memory system designs, and the reuse distance metric we introduce can be useful within the hybrid main memory system context – an avenue we leave for future work.

4. L2ARC does *not* buffer write requests in the SSD and, hence, we do not provide the flow of write accesses in Fig. 2.

blocks in the SSD, L2ARC mainly improves the performance of *Read After Read* (RAR) accesses with a reuse distance larger than the DRAM size (i.e., non-popular blocks), which has a relatively small impact on overall performance.

Dell EMC *Fully Automated Storage Tiering* (FAST) cache [3] is an enterprise approach that employs a two-level I/O cache in storage products. This scheme sets the write policy of the DRAM cache (called SP⁵ cache) to *Write Back* (WB) to accelerate the performance of write requests and uses battery backups to maintain reliability in the presence of power and system failures. The write policy of the SSD cache is also set to WB. Fig. 2b and Fig. 2c show how FAST handles read and write requests. It employs a very simplistic method to identify hot data blocks where a block with more than three accesses in a recent time interval is identified as hot. Hot blocks are promoted to the SSD cache. There is no specified rule for evicting data blocks from cache. Thus, FAST may evict hot data blocks early from the cache (due to promoting new data blocks), which affects both performance and endurance of the SSD. In this case, promotion of the evicted hot data block again imposes additional write operations on the SSD.

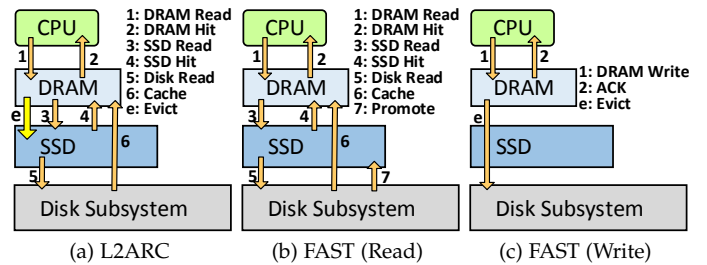


Fig. 2: Architecture of two-level I/O caches employed in enterprise storage systems.

Table 1 summarizes the previously discussed architectures. Among all previous proposals, uCache [37], L2ARC [33], [34], [35], [36], and FAST [3] as two-level caching schemes and ECI-Cache [6] as a hypervisor-based I/O caching scheme are close to our proposal. However, the previous two-level caching schemes neglect cache write policy assignment on different cache levels and hence, fail to provide storage reliability and high endurance for the SSD cache. L2ARC works in the file system level and only improves read performance without considering the endurance of the SSDs. L2ARC fails in detecting popular blocks and does not provide efficient eviction/promotion methods where the cache is managed in a simple FIFO mode. FAST cache keeps write pending data blocks in volatile DRAM cache and employs additional battery backups to provide reliability and hence imposes additional cost. However, such scheme cannot guarantee the reliability of write pending data blocks. In addition, FAST cache employs a simple method to recognize hot data blocks which fails in different types of workloads. uCache keeps both read and write data blocks in DRAM. Thus, a power failure leads to data loss. On the other hand, the DRAM cache becomes polluted due to caching all requests in such small space. In addition, uCache does not consider the popularity of data blocks in the proposed caching scheme. ECI-Cache dynamically assigns different write policies on different VMs caches but does not take into account the impact of cache write policy on cache size estimation and hence, overestimates cache sizes for the VMs that use write policies other than WB. ECI-Cache imposes additional write operations to the SSD cache by updating the cache on each miss operation. In

5. SP: Storage Processor

addition, it is one level: it employs only an SSD, no DRAM. Buffering both read and write requests in the SSD reduces the SSD endurance.

TABLE 1: Summary of previous architectures (WSS: Working Set Size, RI: Reuse Intensity, TRD: Traditional Reuse Distance, URD: Useful Reuse Distance, and POD: Policy Optimized reuse Distance).

Architecture	Industrial	Virtualized	Cache Space Partitioning	Multi Level	Write Policy Management	Endurance Improvement	Reliability Improvement
S-CAVE [10]	✗	✓	WSS	✗ (SSD)	✗	✗	✓
vCS [9]	✗	✓	RI	✗ (SSD)	✗	✓	✓
Centaur [11]	✗	✓	TRD	✗ (SSD)	✗	✗	✗
ECI-Cache [6]	✗	✓	URD	✗ (SSD)	✓	✓	✗
Argon [57]	✗	✗	Acc. Based	✗ (DRAM)	✗	-	✗
Janus [12]	✗	✗	Read Acc. Based	✗ (SSD)	✗	✗	✗
Hystor [7]	✗	✗	Global	✗ (SSD)	✗	✗	✗
ReCA [8]	✗	✗	Global	✗ (SSD)	✓	✓	✗
[22]	✗	✗	Global	✓ (DRAM+SSD)	✗	✓	✗
uCache [37]	✗	✗	Global	✓ (DRAM+SSD)	✗	✗	✗
L2ARC [33], [34], [35], [36]	✓	✗	Global	✓ (DRAM+SSD)	✗	✗	✗
FAST [3]	✓	✗	Global	✓ (DRAM+SSD)	✗	✓	✗
ETICA (Proposed)	✗	✓	POD	✓ (DRAM+SSD)	✓	✓	✓

3 MOTIVATION AND ILLUSTRATIVE EXAMPLE

I/O caches are employed to enhance the performance of storage systems. High-locality data blocks are kept in the cache to serve future requests to them faster and hence reduce the response time significantly. Currently, enterprise storage systems employ high-performance SSDs in the I/O caching layer. SSDs provide more than 500X the performance provided by HDDs (in terms of IOPS for random requests) but SSDs wear out fast due to write operations. Depending on the cache *write policy*, both read and write requests from the application layer may impose write operations on the SSD. Four different write policies can be used on the cache. We elaborate on how 1) requests are supplied by the cache, and 2) performance, reliability, and endurance are affected by different write policies:

- **Write Back (WB)** (with write allocate and no read through⁶) buffers write requests in the I/O cache and writes them back to the storage subsystem when the dirty blocks are evicted. Read requests are also buffered in a WB cache. Each miss from the cache imposes an additional write operation into the cache to store the read block. The WB policy reduces write accesses to the disk subsystem and thus improves both read and write performance. WB policy may lose the write-pending data before they can be written back to the storage subsystem, on power failure. In addition, an SSD WB cache wears out fast due to the large number of write operations induced on the cache.

6. A read through cache buffers read misses in the cache, and hence, serves further read accesses to that address.

- **Write Through (WT)** (with write allocate and no read through) buffers write requests in the I/O cache and at the same time also commits them to the storage subsystem. In addition to writes, WT cache buffers also read requests. WT does *not* improve write performance and only helps read performance. WT policy preserves the reliability of writes since it does *not* buffer any write-pending data (i.e., dirty blocks). WT policy has the same negative impact on the endurance of the SSD as the WB policy.
- **Write Only (WO)** (with write allocate and read through) handles the write requests in a similar way to the WB policy. Read requests are *not* buffered in the WO cache. Hence, this policy improves the performance of only write and *Read After Write* (RAW) operations. WO does not incur additional write operations due to read misses in the cache and, therefore, significantly improves the endurance of an SSD cache.
- **Read Only (RO)** (with no read through) buffers *only* read requests and directs writes to the storage subsystem. RO improves only read performance and preserves the reliability of write requests. RO has a positive impact on SSD cache endurance because it does not expose the I/O write requests to the cache.

We conduct comprehensive experiments to examine the impact of these write policies on the performance and endurance of the SSD cache (in terms of number of write operations on the SSD). The results of these experiments motivate us in selecting the optimized cache policies to provide higher performance, endurance, and reliability in our proposed I/O caching architecture. In these experiments, we study the use of an SSD as a caching layer for the HDD-based storage subsystem with three types of write policies: 1) WB, 2) RO, and 3) WBWO (i.e., WB and WO). WT cache compared to WB (as described earlier) provides similar performance for read accesses and degraded performance for write accesses (due to simultaneous writes in both SSD and HDD devices). Hence, we do not report WT cache performance in our evaluations.

To do so, we perform experiments on a real system. The experimental setup and the running workloads are reported in Table 2. We use an open-source EnhanceIO cache module [70] to implement our I/O caching scheme. The system is warmed up for 30 minutes and then we run the workload for one hour.

TABLE 2: Setup of the motivational experiments.

Hardware						
HDD	6x 4TB SAS 7.2K Seagate HDD in R5 (5+1) configuration Disk Partition size: 200GB					
SSD	4x 2TB Samsung 863a SSD in K10 (2+2) configuration Disk Partition size: 50GB					
DRAM	128GB Samsung DDR4					
CPU	32xIntel(R) Xeon (R) 2.1GHz CPU core					
Software						
OS	CentOS 7 (kernel version: 3.10.327)					
Workloads						
FIO [71] RandRW	Req. Size	Req. Type	Access Pattern	IO Depth	Threads	IO Engine
	8KB	Read/Write (Read: 70%)	Random (distribution: zipf1.1)	16	16	Libaio
FileBench [72]	Req. Size	Req. Type	Access Pattern	Threads	WSS	
Web Server	64KB	Read/Write	Random	100	65GB	
Video Server	256KB	Read	Sequential	48	120GB	
Varmail	64KB	Read/Write	Random	16	62GB	

Fig. 3 shows the results of the experiments, in terms of both performance (measured as *Input/Output Operations Per Second* (IOPS)) and endurance (measured in terms of committed write operations on the SSD). We make four major observations:

- 1) In the FIO-RandRW workload, WB achieves the highest performance. But, WB also imposes the highest number of




Fig. 3: Effect of the cache write policy on the performance of workloads and endurance of the SSD cache (WB: Write Back, RO: Read Only, WBWO: Write Back and Write Only, and IOPS: Input/Output Operations Per Second).

write operations on the SSD compared to other policies. As shown in Fig. 3a, by using WBWO cache, performance decreases by 10.7% while the number of write operations on the SSD also decreases by 32%. The RO cache provides about 78% lower performance compared to the WB cache while it decreases the SSD writes by 82%. This is because the FIO-RandRW workload mainly includes *Read After Write* (RAW) operations (70% read operations), where buffering the write requests in the cache (as in the WB and WBWO policies) significantly improves both read and write performance. WBWO policy does not cache read misses and hence decreases the number of SSD writes by 32% with only 10% performance degradation, compared to the WB cache. On the other hand, RO does not help the performance of RAW operations and, therefore, cannot provide good performance. In addition, RO does not cache I/O write requests and hence, provides no performance improvement for them.

- 2) In the Web Server workload, the WB policy achieves the highest performance while also having the highest number of write operations on the SSD. The difference in performance provided by the three policies is very low: Fig. 3b shows that the WB cache achieves about 15% higher performance compared to WBWO by imposing 98% more writes on the SSD. Compared to RO, WB only marginally improves performance (by 2.7%) at the cost of imposing 11.3% more writes on the SSD. This is because the Web Server workload mainly includes random cold reads (i.e., the first read access to an address) with low locality of reference, and keeping such data blocks in the SSD cache does not help in improving read performance. The WB and WO caches keep read miss data blocks in the SSD cache and impose about 178M and 158M writes on the SSDs, respectively, while WBWO, compared to WB, does *not* cache read misses and achieves almost similar read performance (about 15% lower) with a significantly smaller number of writes on the SSDs (about 98% lower).
- 3) In the Video Server workload, all three policies achieve the same read and write performance with considerable difference in the number of SSD writes. As shown in Fig. 3c, WB, RO, and WBWO achieve 382 IOPS in read and 0 IOPS in write requests (the running workload includes *only* read requests). WB and RO impose about 29M writes on the SSD cache while WBWO does *not* perform any write operations on the SSD cache. This is because the Video Server includes sequential read requests without any locality of reference. Thus, WB and RO perform a large number of unnecessary writes (due to read misses) into the SSD cache without any performance gain. WBWO imposes no writes to the SSD cache and achieves the same

performance as WB and RO. Note that for such a sequential read workload, all requests are supplied from the HDD and the cache hit ratio is equal to 0.

- 4) In the Varmail workload, WB has the highest performance and SSD writes. Both WBWO and RO reduce performance and SSD writes over WB. This is because the RO and WBWO caches fail in serving RAW and RAR accesses, respectively. Fig. 3d shows that both read and write performance of WB are about 41% more than WBWO while WBWO imposes about 40% fewer writes on the SSD. Compared to RO, WB achieves 61% higher performance by imposing 73% more writes into the SSD. The reason is that Varmail includes an equal number of random read and write requests. WBWO does not cache read misses and hence cannot help the RAR requests, but it reduces the number of writes on the SSD. RO does not buffer write requests and thus RAW requests, but it reduces the number of writes on the SSD.

Our main experimental conclusions are:

- 1) In most of the workloads, WB cache achieves the highest performance for read and write requests, however it also imposes the highest number of writes on the SSD. Hence, WB does *not* balance performance and endurance.
- 2) In workloads with a large number of cold reads (such as Web Server and Video Server), WBWO has almost the same performance as WB while it imposes a much smaller number of writes on the SSD cache.
- 3) In workloads with a large number of read requests (such as FIO-RandRW), WBWO cannot provide as good read or write performance as WB, but it results in a smaller number of SSD writes.
- 4) RO cache can be employed in read-intensive workloads with a large number of re-references to greatly extend the SSD lifetime with negligible performance impact compared to WB.
- 5) In workloads such as Video Server and Web Server, cache read misses impose a large number of write operations on the SSD, affecting both endurance and performance.

4 ETICA ARCHITECTURE

In this section, we present the architecture of our proposed two-level I/O cache, ETICA. ETICA has four major characteristics. It: 1) employs both DRAM and SSD in the caching layer of virtualized platforms, 2) assigns effective cache write policies to the two levels of cache, 3) effectively detects the popular data blocks and aims to evict only unpopular blocks from the cache, and 4) allocates efficient cache space for the VMs using the *Policy Optimized reuse Distance* (POD) metric.

Fig. 4 provides an overview of the proposed two-level cache for virtualized platforms. The I/O cache employs two

levels including DRAM in the first and the SSDs in the second level. The information about the content of cache levels is kept in a table called *Map*. The hypervisor receives the I/O requests of the VMs, and routes them to the storage subsystem. The two-level cache resides on the path of requests to the disk subsystem. The cache searches the destination address of the requests in *Map* and finds out whether the corresponding data is located in the 1st level (i.e., DRAM hit) or the 2nd level (i.e., SSD hit). If neither is the case, the request is supplied by the disk subsystem (i.e., cache miss).⁷

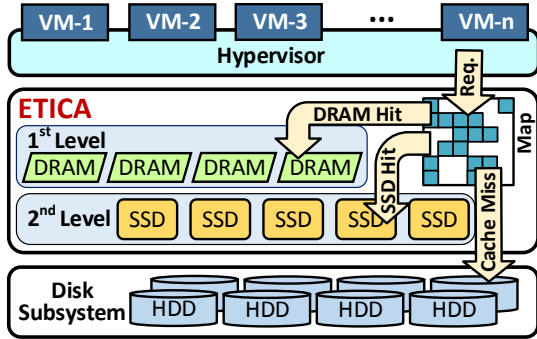


Fig. 4: Architecture of ETICA.

We aim to take advantage of the different merits of both the DRAM technology and SSDs while minimizing the negative characteristics of each. DRAM provides high performance but comes with 1) reliability issues (due to volatility) and 2) high cost. SSDs preserve reliability (due to non-volatility) and have 20X lower cost than DRAM, but they have 1) lower performance and 2) limited endurance (i.e., write cycles). As shown in Fig. 4, DRAM in the first level of the proposed cache mainly improves performance (and endurance) while the SSD in the second level avoids the performance spikes due to the large performance gap between DRAM and the disk subsystem and provides reliability. In Section 4.1, we describe how we assign different write policies at different levels of cache to address the shortcomings and exploit the advantages of each technology. In Section 4.2, we propose popular/unpopular block detection and our promotion/eviction schemes to manage the two-level cache even more efficiently. In Section 4.3, we describe how ETICA estimates efficient cache sizes for different VMs.

4.1 Write Policy Management

We use DRAM in the first level of our I/O cache architecture to enable high performance. Using such a volatile storage technology increases the risk of data loss under different types of failures, such as power failures. We use SSDs in the second level of our I/O cache architecture. SSDs mainly suffer from endurance, i.e., they can support only a limited number of reliable writes. This section shows how we 1) preserve storage reliability in the presence of DRAM and 2) enhance SSD endurance in our proposed two-level I/O cache by applying effective write policies on both levels. Previously, in Section 3, we described the different types of cache write policies and their impact on performance, reliability, and endurance.

To address the storage reliability issue with using DRAM, we assign the RO policy to the first level of cache (i.e., the DRAM level). In this case, *no* write requests are served (i.e.,

supplied) by DRAM and hence, there is no write-pending request (i.e., dirty block) in DRAM. Write requests are directed immediately to the non-volatile second level (i.e., the SSD level), which is able to protect the write-pending data in case of power outage (it is important to note that in our scheme, we use SSDs in the RAID10 configuration [73], which tolerates the failure of up to two SSDs). In this scheme, the DRAM level is responsible for buffering and serving only read requests. Buffering only read requests in DRAM guarantees that any data block in the DRAM level has a copy in another level (i.e., the SSD level or the disk subsystem). Thus, losing data in DRAM has no negative reliability impact.

To address the endurance issue of the SSDs, we assign the WBWO policy to the second level (i.e., the SSD level). Recall that in Section 3, we observed the significant negative impact of buffering read misses on the endurance of the SSD cache. The WBWO policy buffers only write requests in the cache and does *not* buffer read requests. Buffering only write requests in the SSDs effectively improves the endurance of the SSDs by reducing the number of unnecessary write operations on the SSDs (that would otherwise be needed to buffer read requests). Using WBWO at the SSD level does not degrade performance because read requests are buffered and served at the DRAM level.

In summary, the key advantages of our heterogeneous write policy assignment for two different cache levels are:

- 1) We preserve the reliability of write requests by bypassing DRAM and buffering them only in SSDs (in RAID10 configuration).
- 2) We improve the endurance of the SSDs by buffering only write requests while read misses (which would impose a large number of write operations on the SSDs) are only cached in the DRAM level.
- 3) The DRAM level improves the performance of read requests while the SSD level improves the performance of write requests, compared to HDDs. The read misses from DRAM that happen to be supplied by the SSDs (due to RAW requests) also experience the higher performance of SSDs versus the disk subsystem.

We now provide an example and show how the proposed write policy assignment approach in our two-level cache 1) improves the endurance of SSDs, 2) preserves the reliability of I/O requests, and 3) keeps performance intact. In the example of Fig. 5, we compare a one-level WB SSD cache (Fig. 5a) with our two-level cache where the write policy of the DRAM and SSD levels are RO and WBWO, respectively (Fig. 5b). Each cache level has a capacity of 3 data pages. Note that in this example, no cache size estimation algorithm is employed.

When one-level SSD caching is used, as shown in Fig. 5a, *Req*₁ reads the data of sector 1 and buffers the corresponding data block into the cache (due to read miss) by imposing one write operation on the SSD. Similarly, *Req*₂ and *Req*₃ are read misses that are buffered in the cache and each one imposes an additional write on the SSD. *Req*₄ overwrites the data content of the cache (sector 1 in cache block 1) and *Req*₅ writes the contents of sector 4 into the cache (in cache block 2). *Req*₆ hits in the cache and reads sector 1 from cache block 1, which was previously written by *Req*₄. Similarly, *Req*₇ reads data from cache as it is a read hit. We observe that, in the one-level WB SSD cache, the sample workload imposes 5 write operations to the SSD and obtains 2 read hits.

Fig. 5b shows how the two-level cache design of ETICA (with RO DRAM and WBWO SSD levels) supplies the I/O requests with a much smaller number of write operations on the SSD. In the two-level cache, *Req*₁, *Req*₂, and *Req*₃ buffer data blocks into the DRAM level (due to read misses) without imposing any write operations on the SSD. *Req*₄ overwrites

⁷ Each cache level works based on set-associative mapping scheme where each set size is 512 blocks (the set size is configurable to other sizes, but to have a fair comparison, in our experiments the associativity configuration in ETICA is done the same as ECI-Cache).

REQ ID	1	2	3	4	5	6	7
TYPE	Read	Read	Read	Write	Write	Read	Read
SECTOR	1	2	3	1	4	1	4
STATE	Miss	Miss	Miss	Hit	Miss	Hit	Hit

SSD	CB	S	CB	S	CB	S	CB	S	CB	S	CB	S
1	1	1	1	1	1	1	1	1	1	1	1	1
2			2	2	2	2	2	2	2	2	2	2
3			3	3	3	3	3	3	3	3	3	3

SSD Writes: 5, Read Hits: 2

(a) One-level SSD Cache

REQ ID	1	2	3	4	5	6	7
TYPE	Read	Read	Read	Write	Write	Read	Read
SECTOR	1	2	3	1	4	1	4
STATE	Miss	Miss	Miss	Hit	Miss	Hit	Hit

DRAM	CB	S	CB	S	CB	S	CB	S	CB	S	CB	S
1	1	1	1	1	1	1	1	1	1	1	1	1
2			2	2	2	2	2	2	2	2	2	2
3			3	3	3	3	3	3	3	3	3	3

SSD	CB	S	CB	S	CB	S	CB	S	CB	S	CB	S
1			1	1	1	1	1	1	1	1	1	1
2			2	2	2	2	2	2	2	2	2	2
3			3	3	3	3	3	3	3	3	3	3

SSD Writes: 2, Read Hits: 2

(b) Proposed two-level DRAM+SSD cache

Fig. 5: Comparison of a) one-level and b) the proposed two-level cache in terms of endurance, reliability, and performance (CB: Cache Block and S: Sector).

the cache content (sector 1): the write request is served by the SSD and the existing data block in DRAM is marked as invalid. Req_5 is a write request that bypasses DRAM and writes data directly into the SSD. Finally, Req_6 and Req_7 are read requests that hit in the cache and read the data blocks from the SSD level. Unlike the one-level SSD cache, our two-level cache design imposes only 2 write operations on the SSD (60% fewer than one-level) while providing the same hit ratio as one-level does. In summary, compared to the one-level SSD cache, our two-level DRAM+SSD cache with intelligent cache write policy assignment reduces the number of write operations on the SSD while preserving the reliability of write requests and achieving similar performance.⁸

4.2 Promotion/Eviction Optimization

In this section, we propose another approach for improving both endurance and performance in our two-level I/O cache, in which we mainly focus on improving performance (in terms of hit ratio) with the minimum number of write operations into the SSD. The proposed approach is mainly adopted from *pull mode caches* (i.e., *non-datapath caches*), where the content of the SSD cache is not updated during handling of miss accesses (as opposed to *push mode caches* or *datapath caches*) [74]. Updating cache blocks on each cache miss (as *push mode caches* such as LRU⁹ and LFU¹⁰ caches do) leads to promoting new blocks to the cache which could easily be less popular than the evicted ones. Instead, periodically, we detect *unpopular* data blocks and evict them from the SSD and promote *popular* ones. Thus, we minimize the probability of evicting *popular* blocks from the cache, thereby avoiding performance (due to cache miss) and endurance (due to promoting new data blocks into the cache) overheads associated with the churn caused by the pollution due to unpopular blocks. To do so, we first detect *popular* and *unpopular* data blocks and then decide *when* to promote and evict them to/from the SSD level.

We aim to detect *popular* and *unpopular* data blocks via online characterization of the running workload on the VMs. To this end, in specific time intervals (which are determined

8. Performance of our scheme would be higher when requests hit in the DRAM level.

9. Least Recently Used

10. Least Frequently Used

based on the number of requests, e.g., after observing N I/O requests), we analyze the characteristics of the running workloads based on the collected information such as 1) request type, 2) number of accesses, and 3) POD of data blocks. In Section 4.2.1, we describe how we detect *popular* and *unpopular* data blocks.

Fig. 6 shows the flow of I/O requests in our proposed two-level I/O cache for both read and write accesses. As shown in this figure, we employ two queues in our scheme: 1) *promotion queue* in the disk subsystem level and 2) *eviction queue* in the SSD cache level. The data blocks in the disk subsystem that are recognized as *popular* are pushed into the promotion queue while the *unpopular* ones in the SSD cache are inserted into the eviction queue. Periodically, we evict the blocks in the eviction queue from the SSD cache to the disk subsystem and promote the blocks in the promotion queue into the SSD cache (only when there is free space in SSD). This approach greatly increases the likelihood that 1) data blocks are evicted from the SSD when they become *unpopular*, 2) only *popular* data blocks are buffered in the SSD cache, and 3) we do not replace a data block in the SSD cache with a less popular one. Hence, the proposal improves both performance and SSD endurance.

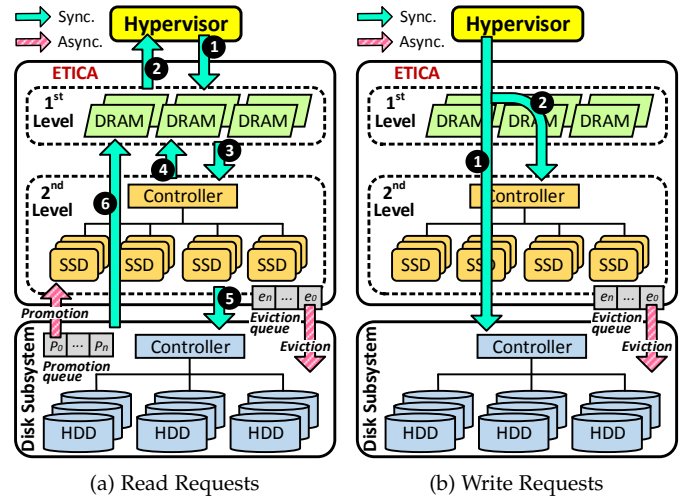


Fig. 6: Examples showing how ETICA handles requests.

We now elaborate on how our proposed promotion and eviction approach in our two-level I/O cache handles the I/O requests. Fig. 6a and Fig. 6b provide the detailed flow of read and write requests, respectively.

4.2.0.1 **Read:** As shown in Fig. 6a, the read request is first received in the DRAM level (1). If the corresponding data block is found in the DRAM cache (i.e., DRAM hit), the request is served via the DRAM level (2). Otherwise, the request is sent to the SSD in the second level (3). In the case of an SSD hit, we first promote the data block to the DRAM level (4) and then serve the request. Otherwise, in case of an SSD miss (5) we read the data from the disk subsystem and promote it directly to the DRAM level *without* buffering in the SSD level (6). Recall that the write policy management scheme presented in Section 4.1 sets the policy of the SSD level to WBWO, where no read miss is buffered in the SSD. However, any data block that is detected as *popular* in the disk subsystem (listed in the promotion queue) is promoted to the SSD level to accelerate future accesses. Furthermore, *unpopular* data blocks in the SSD are evicted through the eviction queue.

4.2.0.2 **Write:** Fig. 6b shows how we handle write requests in the proposed two-level I/O cache. Write requests

(in case of both hit or miss) bypass the DRAM (since the write policy of the first level is set to RO). In case of an SSD miss (①), the data block is directly written to the disk subsystem. Otherwise, in the case of an SSD hit where the corresponding data block is located in the second level (as a *popular* block), the data is written to the SSD level (②). Note that in case of data block update in any level (SSD and disk subsystem), we invalidate the corresponding data block in upper level (DRAM and SSD, respectively). The operation of the eviction queue in handling writes is the same as that of the eviction queue in handling read requests.

To summarize, the proposed eviction/promotion scheme provides the following key benefits:

- 1) Performance improvement by keeping *popular* data blocks in the SSD cache as long as possible.
- 2) SSD endurance improvement (reduced number of SSD cache updates) by periodically promoting *only* popular data blocks into the SSD cache.

4.2.1 Popular and Unpopular Block Detection

In this section, we provide our approach to detecting *popular* and *unpopular* data blocks. Our proposal decides the popularity of data blocks based on two key parameters: 1) POD and 2) frequency of accesses. Hence, we consider both *spatial* and *temporal* locality. We update the $popularity(B_i)$ parameter for each access to the i^{th} data block B_i and push the data block into the eviction or promotion queues based on its relative popularity. The least popular 5% of the blocks in the SSD (i.e., the *unpopular ones*) are put into the eviction queue. Similarly, the most popular 5% of the data blocks in disk subsystem (i.e., HDD) are pushed into the promotion queue. Eq. 1 shows how we calculate $popularity(B_i)$:

$$popularity(B_i) = \sum_{t=1}^{numAcc} e^{-\frac{POD(i,t)}{cacheSize}} \quad (1)$$

Where $POD(i, t)$ is the POD of B_i in the t^{th} access to that data block,¹¹ $numAcc$ is the number of accesses to B_i , and $cacheSize$ is the allocated cache size to the VM. Popularity of each data block is updated periodically in specified time intervals. To keep the popularity information of each data block, a space of 8B is required per data page (4KB). The memory overhead is less than 0.15% (e.g., we need an 8MB space to store the popularity of one million data blocks. This information is kept in SSD). Since popularity is computed asynchronously, its computation does not cause any overhead to the servicing of I/O requests. Fig. 7 shows the impact of POD on the $popularity$ of each access for different cache sizes. It can be seen that calculating $popularity(B_i)$ based on Eq. 1 provides the following key ideas:

- 1) A larger POD leads to smaller $popularity$. Note that when the POD of a block is close to the cache size, the frequency of access to that block is low. Thus, that block likely has no significant impact on the total cache hit ratio.
- 2) When the POD of a block is larger than the cache size, and since this block will be missed from cache, we set low $popularity$ to that block according to Eq. 1. On the other hand, we set high $popularity$ for a block with a POD smaller than the cache size.
- 3) The frequency of the accesses is considered in the popularity calculation. In Eq. 1, we estimate the total popularity of a single block based on all accesses to that block.

The proposed popularity detection scheme is scalable, which is compatible with the removal or addition of VMs into




Fig. 7: Popularity as a function of POD and cache size.

the system. When we add a new VM to the system (or when an existing VM comes online), based on the I/O accesses of that VM, ETICA first assigns efficient cache space in both DRAM and SSD levels and then estimates the popularity of that VM's accesses (we calculate popularity of each VM's accesses independently from other VMs). In contrast, when a VM is removed from the system (or when it goes offline), we reclaim the allocated cache space and thus the popularity detection process for this VM is also stopped. Furthermore, since in specific intervals, ETICA recalculates the cache size of the VMs, in case of extending or shrinking the total SSD and DRAM cache spaces, ETICA is able to reconfigure the newly allocated cache sizes to the VMs based on new SSD and DRAM sizes.

4.3 Cache Size Estimation

In this section, we propose our two-level cache size estimation approach that aims to 1) maximize the overall performance of the co-running VMs and 2) minimize the overall cost (in terms of allocated cache size). We effectively partition the space of both levels (DRAM and SSD) between VMs based on their demand. In Section 4.3.1, we propose the metric of *Policy Optimized reuse Distance* (POD) and demonstrate how the proposed metric effectively reduces the allocated cache size to the workloads (compared to the previous state-of-the-art scheme, URD) by considering both 1) *request type* and 2) *cache write policy* in the reuse distance calculation, and thereby resulting in reduced cost. Then in Section 4.3.2 we show how ETICA assigns efficient cache size for the VMs to achieve the most efficient performance per cost.

4.3.1 Policy Optimized reuse Distance (POD)

Traditional Reuse Distance (TRD) [58], [75], [76], [77], [78], [79], [80], [81], [82], [83], [84], [85] calculates the distance of the requests *only* based on the address and access order of the requests. *Useful Reuse Distance* (URD) [6] improves upon TRD by considering *request type* in reuse distance calculation: it considers *only* the distance of *Read After Read* (RAR) and *Read After Write* (RAW) accesses, enabling a smaller cache with the same or with better performance [6]. However, existing schemes that are employed in cache size estimation neglect the impact of *cache write policy* on reuse distance calculation. In this section, we first provide examples to show the effect of the *cache write policy* on reuse distance calculation of the workloads. Then, we propose a novel metric, *Policy Optimized reuse Distance* (POD), which allocates much smaller cache space compared to TRD and URD while preserving performance (in terms of hit ratio). The POD metric considers both *request type* and *cache write policy* in reuse distance calculation, and hence, does *not* reserve cache blocks for the requests which would *not* be served by the cache (i.e., read accesses in a WBWO and write accesses in a RO cache, respectively). We provide sample workloads and elaborate on how URD (which

11. In Section 4.3.1 we elaborate how to calculate POD.

considers only the request type without considering the cache write policy) and POD (which considers *both* request type and the cache write policy) estimate cache size for caches that use 1) WBWO and 2) RO write policies.

Fig. 8 compares the cache size estimation provided by URD and POD in a WBWO cache. For the given sample workload in Fig. 8a, URD detects the *maximum* reuse distance between Req_1 and Req_6 (due to their RAR accesses to the same sector, i.e., sector 1). Hence, the maximum URD is equal to 4, and 5 blocks of cache are allocated to the workload. Fig. 8a shows that the workload uses only two blocks of the allocated cache space for write requests (write accesses to sector 4 and sector 5 by Req_4 and Req_5 , respectively). Since data blocks of *read misses* are *not* buffered in a WBWO cache, three blocks of allocated cache remain *unused* while they are reserved for this workload. At the end of the workload, Req_7 reads the buffered data block from the cache: it is a read hit due to the RAW operation. We observe that a cache with WBWO policy helps only the write and RAW¹² requests without providing any performance improvement for read and RAR (i.e., Read After Cold¹³ Read) accesses.

The POD metric considers both *request type* (similar to URD) and *cache write policy* in reuse distance and cache size estimation. Fig. 8b shows how POD estimates cache size for the sample workload. POD estimates the *maximum* reuse distance based on Req_4 and Req_7 due to their RAW access to sector 4. In this case, the POD of the workload is equal to 1. Crucially, note that the read access of Req_6 is *not* considered in POD calculation in the WBWO cache. This is because such a request would *not* be served by the cache. Therefore, only two blocks of cache are allocated to the workload. Req_4 and Req_5 commit write requests into the cache (due to their accesses to sectors 4 and 5) and Req_7 (to sector 4) is supplied from the cache. In summary, for the given sample workload, we observe that in a WBWO cache URD allocates 5 blocks while POD allocates only 2 blocks of cache and achieves the same I/O performance (for both read and write requests) as URD.

REQ	TYP	SEC	Cache			S
1	R	1				M
2	R	2				M
3	R	3				M
4	W	4	4			H
5	W	5	4	5		H
6	R	1	4	5		M
7	R	4	4	5		H

(a)

REQ	TYP	SEC	Cache	S
1	R	1		M
2	R	2		M
3	R	3		M
4	W	4	4	H
5	W	5	4	H
6	R	1	4	M
7	R	4	4	H

(b)

Fig. 8: Comparison of cache size allocation in a WBWO cache (a) without and (b) with considering cache write policy (REQ: Request, TYP: Type, SEC: Sector, W: Write, R: Read, S: State, H: Hit, M: Miss).

Fig. 9 shows how URD and POD work in a RO cache, by comparing the cache allocation and I/O performance of these two schemes. For the given sample workload in Fig. 9a, URD detects the maximum reuse distance for the RAW access by Req_1 and Req_7 to sector 1 of the disk. In this case, the URD of the workload is equal to 4 and hence, five blocks of cache are allocated to the workload (note that read access of Req_6 to sector 3 is duplicated with Req_3 and both of them

are supplied by the same cache block). Fig. 9a shows that the write accesses of the workload (Req_1 , Req_4 , and Req_5) are bypassed from the cache and are supplied by the disk subsystem. Only the *read misses* (Req_2 and Req_3) are buffered in cache and RAR accesses such as Req_6 are served while RAW accesses such as Req_7 cannot be served by the RO cache. Only three blocks of allocated cache are actually used and the remaining two blocks are unused but they are reserved by the workload. This is because the RO cache only buffers reads and serves RAR accesses. In this case, reserving cache blocks for write accesses (as URD does) has no performance benefit while imposing additional cost (because it allocates larger cache space with unused cache blocks).

In Fig. 9b, we show how POD considers the *cache write policy* in reuse distance calculation and results in 1) a much smaller cache space allocation compared to URD and 2) the same I/O performance as URD. In the RO cache, POD detects the maximum reuse distance of the workload for the RAR accesses to sector 3 by Req_3 and Req_6 . Since no write request can be supplied by RO cache, write requests are *not* considered in the POD calculation of such a cache. Hence, POD is equal to 0 (Req_4 and Req_5 are not considered since write accesses have no impact on the RO cache operation) and *only one block* of cache is allocated to the workload. As shown in Fig. 9b, write accesses (Req_1 , Req_4 , and Req_5) are not served by the cache. A read miss by Req_2 promotes the data block of sector 2 into the cache. Then, Req_3 updates the cache by promoting the data block of sector 3 (note that in this workload there is no future read access to sector 2 and POD enables updating the data block by Req_3). Req_6 hits in the cache and finally, Req_7 promotes sector 1 into the cache. We observe that in an RO cache, POD calculates the maximum reuse distance based on RAR operations. Compared to URD, POD assigns a much smaller cache space to the workload (only 1 block versus 5 blocks) while achieving exactly the same cache hit ratio.

REQ	TYP	SEC	Cache			S
1	W	1				M
2	R	2	2			M
3	R	3	2	3		M
4	W	4	2	3		M
5	W	5	2	3		M
6	R	3	2	3		H
7	R	1	2	3	1	M

(a)

REQ	TYP	SEC	Cache	S
1	W	1		M
2	R	2	2	M
3	R	3	3	M
4	W	4	3	M
5	W	5	3	M
6	R	3	3	H
7	R	1	1	M

(b)

Fig. 9: Comparison of cache size allocation in RO cache (a) without and (b) with considering cache write policy (REQ: Request, TYP: Type, SEC: Sector, W: Write, R: Read, S: State, H: Hit, and M: Miss).

To summarize, the key ideas behind POD are as follows:

- 1) POD considers both 1) *request type* and 2) the specific *cache write policy* in the reuse distance calculation.
- 2) In a WBWO cache, the reuse distance calculation of POD is based only on the *maximum* distance of RAW accesses (also including RARAWs). This is because no RAR (i.e., Read After Cold Read) request is supplied by a WBWO cache (since the cache does *not* buffer read misses). Hence, POD does *not* reserve cache blocks for read miss requests (note that the read hits are due to RAW accesses where the read cache block is previously buffered by a write request). Thus, POD always estimates a smaller cache size compared to URD, while achieving the *same* hit ratio and

12. Note that WBWO cache supplies *only* RAR accesses where the first read was RAW (i.e., RARAW). We also consider such accesses as RAW and calculate POD for the *maximum* distance of such accesses.

13. The first accesses to an address.

I/O performance (for both read and write operations in a WBWO cache).

- 3) In an RO cache, POD considers the *maximum* distance of only RAR accesses in reuse distance calculation. Since the cache does *not* buffer write requests, no cache block is reserved for write operations by POD. In this case, POD results in a smaller cache size than URD, while providing the same hit ratio and I/O performance.
- 4) In a WB cache, URD and POD work similarly. Both RAR and RAW accesses are supplied by the WB cache. In this case, both URD and POD estimate the cache size based on the *maximum* distance of either RAR and RAW accesses.

4.3.2 Allocating Efficient Cache Size for VMs

We now elaborate on how we allocate an efficient cache size (in both DRAM and SSD levels) for each running VM. ETICA estimates and assigns cache sizes to the VMs by using the POD metric. ETICA aims to achieve the maximum performance (in terms of hit ratio) with the minimum cost (in terms of cache size) as opposed to previous architectures which mainly aim to optimize performance with the maximum available cache sizes [6], [11]. The steps of cache size assignment by ETICA are as follows:

- 1) Periodically, ETICA calculates the POD of the running VMs for both the RO and WBWO policies used in the DRAM and SSD levels, respectively. Recall from Section 4.3.1 that POD does *not* reserve cache blocks for writes in RO and reads in WBWO caches and hence estimates different cache sizes for different write policies.
- 2) Based on the calculated POD, we estimate two efficient cache sizes for the VM in both the DRAM and the SSD levels. Eq. 2 shows how we calculate the cache sizes:

$$\begin{aligned} Cache_{DRAM} &= POD_{DRAM} \times blockSize_{DRAM} \\ Cache_{SSD} &= POD_{SSD} \times blockSize_{SSD} \end{aligned} \quad (2)$$

Where POD_{DRAM} and POD_{SSD} are the calculated POD metrics of the running VM in DRAM and SSD levels and $blockSize_{DRAM}$ and $blockSize_{SSD}$ are the cache block size in DRAM and SSD, respectively.

- 3) We check whether the total estimated cache sizes for all VMs fit in the available physical DRAM and SSD space.
- 4) If the sum of estimated cache sizes exceeds the existing physical DRAM or SSD space, we reduce the estimated size of each VM to maximize the *Performance-Per-Cost* (PPC) factor in Eq. 3.

$$PPC = \sum_{i=1}^{numVMs} \frac{H(VM_i, c_i)}{c_i} \quad (3)$$

Where $numVMs$ is the number of running VMs and $H(VM_i, c_i)$ is the hit ratio achieved by VM_i by allocating it a cache size equal to c_i .

Since the running VMs are weighted identically, we aim to maximize the aggregate *PPC* as proposed in Eq. 3. Assigning cache sizes based on Eq. 3 guarantees the maximum overall performance with minimum cost. To do so, ETICA provides POD-based *Miss Ratio Curves* (MRCs) for each running VM and finds the set of $\{c_0, \dots, c_{N-1}\}$ which provides the maximum *PPC*.

5 EXPERIMENTAL RESULTS

In this section, we evaluate the effectiveness of our proposed two-level I/O caching scheme, ETICA, in terms of cache size, performance, and endurance.

5.1 Experimental Setup

We conduct comprehensive experiments on a real system including a 4U rack mount SuperMicro server with 1) 6x Seagate 4TB SAS 7.2K HDDs in RAID5 (5+1) configuration at the disk subsystem level, 2) 4x Samsung 2TB 863a SATA SSDs configured as RAID10 (2+2) at the SSD cache level, 3) 128GB Samsung DDR4 DRAM where 100GB is used at the DRAM cache level, 4) LSI9361i PCIe RAID controller to apply the RAID configuration, 5) 32x Intel(R) Xeon (R) 2.1GHz CPU cores, and 6) a SuperMicro X10i motherboard. We integrated ETICA with QEMU as the open-source hypervisor running on the CentOS7 operating system (kernel version: 3.10.327.36.3).¹⁴ We implement POD by modifying the source code of PARDA [76].

We run VMs with the *CentOS* and *Ubuntu* operating systems on the hypervisor. Each VM is configured with 1) 50GB disk space, 2) 1GB memory, and 3) two virtual CPUs. ETICA estimates and allocates the I/O cache size for each VM based on the I/O patterns of workloads running on the VMs. We run MSR traces from SNIA [54], [55] in the device layer, including more than ten real application workloads with different I/O characteristics, such as *hm_1* (hardware monitoring), *mds_0* (media server), *mds_1*, *src2_0* (source control), *src1_2*, *stg_1* (web staging), *ts_0* (terminal server), *wdev_0* (test web server), *web_3* (web/SQL server), *rsrch_0* (research projects), *usr_0* (user home directories), and *proj_0* (project directories). To show the effectiveness of ETICA, we compare it with the latest state-of-the-art I/O caching scheme in virtualized platforms, ECI-Cache [6]. To do so, we implemented this scheme in our real testing platform and evaluated it with the same experiments. The evaluation of ECI-Cache and ETICA are performed via the default configuration of the device layer with enabled *request merge* option for a 128-entry device queue size and disabled buffer cache.¹⁵

5.2 Cache Size Improvement

In this section, we empirically show that ETICA reduces the allocated cache space to each VM. Our experiments show how ETICA estimates a smaller cache size compared to the previous state-of-the-art cache space partitioning scheme, ECI-Cache, while preserving the performance (i.e., latency) of the running VMs (Section 5.3). ETICA employs the POD metric in cache size estimation, which effectively allocates smaller cache sizes to the VMs by taking into account both *request type* and *cache write policy* (as we showed in Section 4.3.1). The state-of-the-art scheme, ECI-Cache, works based on URD; we have shown that this metric ignores the impact of cache write policy and only considers the request type in reuse distance estimation and thus overestimates the cache size of the VMs with RO and WBWO policies (see Section 4.3.1). We run 12 VMs with different types of workloads on the hypervisor and measure the cache size estimation of ETICA and ECI-Cache. To do so, we set different cache write policies to the VMs: 1) WB, 2) WBWO, and 3) RO. We estimate the cache size of the VMs in predefined intervals (after observing 10,000 I/O requests) by POD and URD which are used by ETICA and ECI-Cache, respectively.

Fig. 10 shows the VM cache size estimations of POD and URD for different cache write policies. Unlike POD, URD works independently of the cache write policy: that is, URD works *exactly* the same for the caches with WB, WBWO, and

¹⁴. We disable the buffer cache and set the configuration of the block layer to the default mode with *merge* option enabled with a 128-entry device queue size.

¹⁵. Note that ETICA and ECI-Cache has no management on the buffer cache and they only perform in the device layer.

RO policies, and hence, we show the URD results with one single line. To make the differences of cache sizes estimated by the proposed and previous schemes more visible in the figures, we show only the first 20 intervals of the workloads.

¹⁶ In addition, Fig. 11 compares the *average cache sizes* allocated by ECI-Cache (URD) and ETICA (POD) for the VMs over their entire runs. We make five main observations:

- 1) The URD metric allocates the *same* size for the caches with different write policies (i.e., RO, WBWO, and WB). This is because URD *ignores* the cache write policy and *only* considers request type in reuse distance calculation. As we have clearly shown, this scheme allocates cache blocks for requests that would *not* be buffered in cache. For instance, a WBWO cache does not buffer read requests, but URD reserves cache blocks for such requests in its reuse distance calculation. Similarly, URD allocates cache blocks of an RO cache for write requests that bypass the RO cache.
- 2) The amount of cache space allocated by POD in both RO and WBWO policies is considerably smaller (i.e., by 51.7%, on average) than that allocated by URD. This is due to the fact that in an RO cache, POD considers only RAR accesses in the reuse distance calculation and does *not* allocate cache blocks for writes. In a WBWO cache, POD considers only the reuse distance of RAW accesses and allocates no cache block for cold reads (i.e., read misses). In contrast, URD considers both RAR and RAW accesses *without* taking into account the effect of cache write policy on what gets cached and what does not. Unlike POD, the URD metric allocates cache blocks for both writes and reads, which would *not* be buffered in RO and WBWO caches, respectively.
- 3) In the intervals where RAR accesses are involved in cache size estimation (e.g., intervals 10 to 15 in hm_1, 3 to 10

16. We run the experiments for more than 500 intervals (until the workload finishes) and observe very similar behavior across the entire execution.

in proj_0, and 3 to 6 in rsrch_0), URD and POD for the RO cache (POD(RO)) have the same behavior but POD estimates smaller cache sizes (because it does *not* allocate cache blocks for write requests). Similarly, in the intervals where RAW accesses are involved (e.g., intervals 1 to 3 in web_3, 0 to 8 in ts_0, and all intervals of wdev_0), URD and POD for the WBWO cache (POD(WBWO)) have the same pattern but POD estimates a much smaller cache size because it does *not* allocate blocks for cold reads in the WBWO cache.

- 4) In the cache with the WB policy, POD and URD estimate the same cache size. This is because WB cache buffers *both* read and write requests and hence, both schemes allocate cache blocks for both types of requests.
- 5) The POD calculation imposes up to 0.83% performance overhead, which is similar to the overhead imposed by ECI-Cache due to calculating URD (not shown in the figures). We observe that both POD and URD calculations have the same performance overhead on the running VMs.

We conclude that POD provides better, more efficient size estimation for caches that employ RO and WBWO policies, thereby reducing waste of space in an I/O cache.

5.3 I/O Latency Improvement

In this section, we empirically show that ETICA improves the overall performance (in terms of I/O latency) of the running VMs. In the experiments of this section, we run 12 VMs with different workloads on the hypervisor integrated with ETICA and examine the performance in terms of I/O latency of the running VMs. We evaluate the system performance using ETICA under two conditions: 1) without applying the promotion/eviction scheme (denoted as ETICA-NPE) and 2) with applying our proposed promotion/eviction scheme (denoted as ETICA-Full). We perform the same experiments on the ECI-Cache architecture for performance comparison.




Fig. 10: Allocated cache sizes by ETICA and ECI-Cache (URD: ECI-Cache, POD(RO): ETICA at the DRAM level, and POD(WBWO): ETICA at the SSD level).




Fig. 11: Average cache size allocated by ECI-Cache and ETICA (URD: ECI-Cache, POD(RO): ETICA in DRAM level, and POD(WBWO): ETICA in SSD level).

We compare the performance of these two schemes by assuming similar total cache spaces, where the total space of SSD+DRAM in ETICA is *equal* to the space of SSD in ECI-Cache. We recalculate the cache sizes after observing 10,000 I/O requests and we also promote/evict data blocks after observing 1,000 I/O requests (in Section 5.6 we evaluate the impact of promotion/eviction intervals on both the performance and endurance improvement of ETICA).

Fig. 12 and Fig. 13 show the results of the experiments, comparing the average latency and total hit ratio of the running VMs with ETICA and ECI-Cache. We make five observations:

- 1) ETICA, compared to ECI-Cache, improves performance by up to 64% (45%, on average). This is mainly due to the improvement in the latency of read requests. Even without applying our proposed promotion/eviction scheme, ETICA improves the I/O performance by 27%, on average.
- 2) Performance enhancement by ETICA is due to 1) hit ratio improvement and 2) DRAM high I/O performance. As shown in Fig. 13, ETICA provides 30% higher hit ratio compared to ECI-Cache. In this case, allocating DRAM to ECI-Cache cannot close the performance gap between ETICA and ECI-Cache.
- 3) In VM3 running the mds_1 workload, compared to ECI-Cache, ETICA-NPE and ETICA-Full achieve only 1% and 13% performance improvement, respectively. This is due to the sequential read access pattern of the running workload with low locality of reference. As we observe in the experiments of Section 5.2, both ECI-Cache and ETICA assign almost zero cache space for this workload in the first 10 intervals. Using a cache for such workloads with sequential (streaming) read accesses with low locality is unlikely to provide significant performance improvement since most requests would be supplied by the disk subsystem.
- 4) The usr_0 workload mainly consists of write requests, which are supplied by the SSD cache in both ETICA and ECI-Cache schemes. ETICA-Full, compared to ECI-Cache, improves the I/O performance of this workload by 45% since it buffers RAW requests at the DRAM level and does not evict popular written data blocks from the SSD (ETICA-NPE improves the performance of this workload by 27%).
- 5) The running workload in VM1, hm_1, mainly includes random read accesses with high locality of reference. ETICA obtains a high hit rate from DRAM (about 91.6%) for this workload and improves I/O performance by 42%. ECI-Cache achieves a similar hit rate in the SSD, but the higher latency of the SSD compared to DRAM leads to lower performance by ECI-Cache.
- 6) ETICA achieves the highest performance improvement

for the src2_0, ts_0, and wdev_0 workloads (53%, 58%, and 64%, respectively, over ECI-Cache). These workloads include a small number of write operations with a large number of RAW (and also RARAW) re-references, which leads to promoting a small number of written data blocks from the SSD to the high-performance DRAM and serving further accesses to them with much lower latency than ECI-Cache.




Fig. 12: Average latency of VMs with ETICA and ECI-Cache (ETICA-NPE: without promotion/eviction and ETICA-Full: with promotion/eviction).




Fig. 13: Total hit ratio of VMs with ETICA and ECI-Cache.

We conclude that ETICA provides 45% higher I/O performance than ECI-Cache. This improvement is mainly due to 1) the performance improvement of read requests (especially the requests that are supplied by the DRAM level) and 2) not evicting popular data blocks from the SSD level.

5.4 SSD Endurance Improvement

In this section, we evaluate the SSD endurance improvement of ETICA over ECI-Cache. We measure endurance *with the number of writes into the SSD cache* as our metric (as discussed and used by prior works [6], [86], [87], [88], [89]). Fig. 14 compares the endurance of SSD using ETICA and ECI-Cache (in terms of the number of writes performed into the SSD cache). A lower number of writes indicates higher endurance. We make two major observations:

- 1) ETICA reduces the number of writes into the SSD by 33.8% in average. The maximum endurance improvement (about 95%) is achieved in the web_3 workload, mainly due to not buffering read misses at the SSD level. This workload consists mainly of read accesses with a small number of write operations (i.e., it is read-intensive). Thus, applying the RO policy at the SSD level, as ECI-Cache does, does *not* help endurance since *all* read misses (i.e., cold reads) would be buffered in the cache, imposing a large number of writes into the SSD.
- 2) stg_1, src2_0, and rsrch_0 are write-intensive workloads where both ECI-Cache and ETICA buffer write requests

in the SSD cache. ETICA achieves about 24%, 14%, and 16% endurance improvement over ECI-Cache for these workloads, respectively. This is because ETICA does *not* update the SSD cache on each miss and hence, it reduces the number of cache updates by only promoting popular blocks into the SSD cache.

We conclude that ETICA, compared to ECI-Cache, significantly reduces the number of write operations on the SSD cache (by 33.8%, on average) and thus results in improved SSD endurance and lifetime. Read-intensive workloads impose a large number of writes into the SSD cache and ETICA avoids updating the SSD cache with such read requests. To obtain high performance, ETICA buffers read requests at the DRAM level instead of at the SSD level. Thus, the two-level I/O caching architecture of ETICA improves both performance and endurance.




Fig. 14: Number of write operations performed on the SSD cache by ECI-Cache and ETICA.

5.5 Performance Analysis When Enabling New VMs

In this section, we show how ETICA reallocates cache resources for running VMs and how it affects the performance (in terms of cache hit ratio) when a varied number of VMs is enabled. In these experiments, we set the total available cache space equal to limited number of cache blocks (i.e., 50,000 cache blocks). In this case, when the total demand of running VMs exceeds the available cache space, ETICA reduces the allocated cache space with the minimum performance overhead. First we start the experiments with running one VM and then at the specific time instances (i.e., intervals 2, 6, and 10) we extend the number of running VMs to 4, 16, and 32. During intervals 0 to 2, we run *hm_1* workload on VM0.¹⁷ Then at interval 2, VM1, VM2, and VM3 start to run *proj_0*, *stg_1*, and *usr_0* workloads, respectively. At interval 6 we run the following workloads in VMs 4 to 15: *ts_0*, *wdev_0*, *web_3*, *usr_0*, *mds_0*, *usr_0*, *ts_0*, *wdev_0*, *ts_0*, *hm_1*, *src2_0*, and *ts_0*. At the last interval, VMs 16 to 31 start running workloads: *ts_0*, *wdev_0*, *web_3*, *usr_0*, *mds_0*, *usr_0*, *ts_0*, *wdev_0*, *ts_0*, *wdev_0*, *ts_0*, *hm_1*, *src2_0*, *ts_0*, *ts_0*, *wdev_0*, *web_3*, and *usr_0*. The hardware and software configurations of the running VMs in these experiments are the same as reported configurations in Section 5.1.

Fig. 15 and Fig. 16 show the results of the experiments. In Fig. 15, we present how ETICA recalculates new cache sizes for the running VMs based on their demand and available total cache size (equal to 50,000 cache blocks in the experiments). Fig. 16 shows the average cache hit ratio of the running VMs while we run 1, 4, 16, and 32 online VMs.

It can be seen that in the intervals 1 to 10, by enabling new VMs, ETICA allocates required cache sizes to the VMs (Fig.

17. The running workloads are from SNIA and the detailed information about workload characteristics are available in [54], [55].

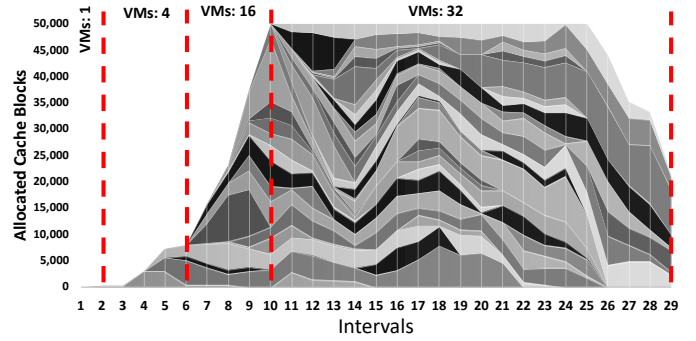


Fig. 15: Cache reallocation by ETICA while enabling different number of VMs (up to 32 VMs, Interval size: 10min).

15). Between intervals 10 to 29 (where we run 32 concurrent VMs), the total required cache space by the VMs is greater than the total cache space. In this case, ETICA reduces the allocated cache sizes, affecting the performance of the VMs (Fig. 16). As shown in Fig. 16, when only VM0 is online, the average hit ratio is equal to 96%, then by enabling VM1, VM2, and VM3, the average hit ratios for these VMs are: 77%, 44%, 29%, and 37%. The performance reduction in VM0 is due to reduced locality of accesses by this VM (in intervals 1-6, the available total cache space is greater than VMs demand). This performance behavior is also observed in another experiment when running only this VM with unlimited cache resources. By adding new VMs (interval 7), ETICA reduces the allocated cache sizes and hence, we experience performance drop in the VMs. In this case, the average hit ratio by VM0 is 28%. Next, ETICA increases the allocated cache to VM0, since other VMs demand is low, and hence, cache hit ratio increases to 64%.

5.6 Analysis of the Promotion/Eviction Intervals

The required promotions to (or evictions from) SSD are conducted when a fixed number of I/O requests are processed by ETICA. Using small intervals enables ETICA to respond faster to changes in the workload characteristics with the cost of more writes in SSD and therefore, decreasing its lifetime. To show the impact of interval size on both performance and endurance of ETICA, we conduct experiments and change the interval length from 100 to 10,000 I/O requests processed by the system. Fig. 17 shows the normalized performance and endurance of ETICA in various promotion/eviction intervals. We make four major observations:

- 1) There is a negligible performance improvement between very small promotion/eviction intervals (less than 1,000) and due to the cost of promotion/evictions, employing very small interval sizes is not efficient.
- 2) The performance difference between various intervals is about 20% in all workloads.
- 3) In benchmarks with steady workload characteristics throughout their runtime, increasing the interval size does not have any side-effect on the performance. Our investigation reveals that this is due to the sparse accesses to the large number of data pages. When the intervals are large, such data pages accumulate a high score and therefore, replace the hot data pages in the SSD.

We conclude that the selected interval size for promotions/evictions in ETICA can efficiently balance the performance and endurance in all workloads.




Fig. 16: Average performance of VMs (in terms of cache hit ratio) while enabling different number of VMs (up to 32 VMs).




Fig. 17: Impact of promotion/eviction intervals on the performance and endurance of ETICA.

5.7 ETICA vs. ECI-Cache

In this section, we summarize how ETICA overcomes the shortcomings of the latest state-of-the-art I/O caching scheme, ECI-Cache, in terms of cache size, performance, and endurance.

Efficient Cache Size Estimation. ETICA employs POD to estimate the cache size, which results in reduced allocated cache sizes to the VMs. Unlike URD, POD considers cache write policy in reuse distance estimation, and thus it does *not* reserve cache blocks for accesses that would not be served by the cache.

Higher Performance. ETICA, compared to ECI-Cache, improves performance in two ways: 1) it employs high-performance DRAM in the two-level I/O cache structure and 2) it promotes (evicts) only popular (unpopular) data blocks in (from) the I/O cache.

Higher SSD Endurance. The proposed per-level write policy management scheme in ETICA effectively reduces the number of write accesses on the SSD level while improving the performance and preserving reliability of write-pending data blocks. In contrast, ECI-Cache buffers both read and write accesses in the SSD cache, which causes more writes into the SSD cache.

6 CONCLUSION

In this paper, we presented ETICA, a new two-level I/O caching scheme for virtualized platforms. ETICA takes advantage of both DRAM and SSD in the I/O caching layer and improves cost, performance, and endurance while preserving the reliability of I/O requests. The write policy of the first caching level (i.e., DRAM) is set to RO (to preserve the storage reliability in the presence of volatile DRAM) while we use the WBWO policy (to improve SSD endurance) in the second caching level (i.e., SSD). DRAM cache enhances the performance of read requests by buffering read misses while write requests are buffered by the SSD in the second level, providing both high performance and high reliability for the write-pending requests. ETICA further improves the

I/O performance of the running workloads by detecting and buffering *popular* data blocks where the data blocks are not evicted from the cache until they become *unpopular*. ETICA improves the endurance of the SSDs by 1) not buffering read requests in the SSD cache, which results in a significantly reduced number of writes into the SSD cache and 2) only promoting *popular* data blocks into the cache and eliminating the update of the SSD in the case of read misses. ETICA improves the cost of the I/O cache by 1) allocating a smaller cache size to each VM and 2) assigning an effective write policy to each cache level and hence, eliminating the need for using high-cost peripherals such as battery backup to maintain the reliability of write-pending data blocks in the case of power or system failures. ETICA employs our new *Policy Optimized reuse Distance* (POD) metric, which considers both 1) *request type* and 2) *cache write policy* in cache size estimation and hence, estimates a much smaller cache size for each VM while preserving the I/O performance of the VMs. The results of our real system experiments show that ETICA provides 45% and 33.8% higher performance and endurance and also 51.7% reduced cache size compared to the best state-of-the-art I/O caching policy [6] in virtualized platforms.

REFERENCES

- [1] V. Soundararajan and J. M. Anderson, "The Impact of Management Operations on the Virtualized Datacenter," in *ACM International Symposium on Computer Architecture*, 2010.
- [2] StorageReview, "StorageReview," <https://www.storagereview.com>, Accessed: Aug. 2019.
- [3] Dell EMC Corp. (2018) Dell EMC Unity: FAST Technology Overview. <https://www.emc.com/collateral/white-papers/h15086-emc-unity-fast-technology-overview.pdf>. [Online]. Available: <https://www.emc.com/collateral/white-papers/h15086-emc-unity-fast-technology-overview.pdf>
- [4] HP. (2018) HPE Smart Array SR SmartCache. [Online]. Available: <https://h20195.www2.hp.com>
- [5] NetApp. (2017) SSD Cache Feature. <https://library.netapp.com>. [Online]. Available: <https://library.netapp.com>
- [6] S. Ahmadian, O. Mutlu, and H. Asadi, "ECI-Cache: A High-Endurance and Cost-Efficient I/O Caching Scheme for Virtualized Platforms," *Proc. ACM Meas. Anal. Comput. Syst. (POMACS)*, vol. 2, no. 1, pp. 9:1–9:34, Apr. 2018. [Online]. Available: <http://doi.acm.org/10.1145/3179412>
- [7] F. Chen, D. A. Koufaty, and X. Zhang, "Hystor: Making the Best Use of Solid State Drives in High Performance Storage Systems," in *International Conference on Supercomputing*, 2011.
- [8] R. Salkhordeh, S. Ebrahimi, and H. Asadi, "ReCA: An Efficient Reconfigurable Cache Architecture for Storage Systems with Online Workload Characterization," *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, vol. 29, no. 7, pp. 1605–1620, July 2018.
- [9] F. Meng, L. Zhou, X. Ma, S. Uttamchandani, and D. Liu, "vCacheShare: Automated Server Flash Cache Space Management in a Virtualization Environment," in *Proceedings of USENIX Annual Technical Conference (USENIX ATC)*, 2014.
- [10] R. Barik, J. Zhao, and V. Sarkar, "S-CAVE: Effective SSD Caching to Improve Virtual Machine Storage Performance," in *Proceedings of Parallel Architectures and Compilation Techniques (PACT)*, 2013.
- [11] R. Koller, A. J. Mashtizadeh, and R. Rangaswami, "Centaur: Host-Side SSD Caching for Storage Performance Control," in *IEEE International Conference on Autonomic Computing (ICAC)*, 2015.

- [12] C. Albrecht, A. Merchant, M. Stokely, M. Waliji, F. Labelle, N. Coehlo, X. Shi, and E. Schrock, "Janus: Optimal Flash Provisioning for Cloud Storage Workloads," in *Proceedings of USENIX Annual Technical Conference (USENIX ATC)*, 2013.
- [13] S. Byan, J. Lentini, A. Madan, L. Pabon, M. Condict, J. Kimmel, S. Kleiman, C. Small, and M. Storer, "Mercury: Host-Side Flash Caching for the Data Center," in *IEEE Symposium on Mass Storage Systems and Technologies (MSST)*, 2012.
- [14] D. Narayanan, E. Thereska, A. Donnelly, S. Elnikety, and A. Rowstron, "Migrating Server Storage to SSDs: Analysis of Tradeoffs," in *ACM European Conference on Computer Systems (EuroSys)*, 2009.
- [15] Q. Xia and W. Xiao, "High-Performance and Endurable Cache Management for Flash-Based Read Caching," *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, vol. 27, no. 12, pp. 3518–3531, 2016.
- [16] J. Matthews, S. Trika, D. Hensgen, R. Coulson, and K. Grimsrud, "Intel® Turbo Memory: Nonvolatile Disk Caches in the Storage Hierarchy of Mainstream Computer Systems," *ACM Transactions on Storage (TOS)*, 2008.
- [17] Y. Kim, A. Gupta, B. Uргаonkar, P. Berman, and A. Sivasubramanian, "HybridStore: A Cost-Efficient, High-Performance Storage System Combining SSDs and HDDs," in *IEEE Annual International Symposium on Modelling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)*, 2011.
- [18] C. Li, P. Shilane, F. Douglass, H. Shim, S. Smaldone, and G. Wallace, "Nitro: A Capacity-Optimized SSD Cache for Primary Storage," in *Proceedings of USENIX Annual Technical Conference (USENIX ATC)*, 2014.
- [19] Y. Klonatos, T. Makatos, M. Marazakis, M. D. Flouris, and A. Bilas, "Azor: Using Two-Level Block Selection to Improve SSD-Based I/O Caches," in *IEEE International Conference on Networking, Architecture and Storage (NAS)*, 2011.
- [20] Q. Xia and W. Xiao, "Flash-Aware High-Performance and Endurable Cache," in *Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS)*, 2015 *IEEE 23rd International Symposium on*. IEEE, 2015, pp. 47–50.
- [21] Y. J. Yu, D. I. Shin, W. Shin, N. Y. Song, J. W. Choi, H. S. Kim, H. Eom, and H. Y. Yeom, "Optimizing the Block I/O Subsystem for Fast Storage Devices," *ACM Trans. Comput. Syst. (TOCS)*, vol. 32, no. 2, pp. 6:1–6:48, Jun. 2014. [Online]. Available: <http://doi.acm.org/10.1145/2619092>
- [22] R. Salkhordeh, M. Hadizadeh, and H. Asadi, "An Efficient Hybrid I/O Caching Architecture Using Heterogeneous SSDs," *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, pp. 1–1, 2018.
- [23] S. Ahmadian, R. Salkhordeh, and H. Asadi, "LBICA: A Load Balancer for I/O Cache Architectures," in *to appear in Design, Automation Test in Europe Conference Exhibition (DATE)*, March 2019.
- [24] J. Ousterhout, A. Gopalan, A. Gupta, A. Kejriwal, C. Lee, B. Montazeri, D. Ongaro, S. J. Park, H. Qin, M. Rosenblum, S. Rumble, R. Stutsman, and S. Yang, "The RAMCloud Storage System," *ACM Trans. Comput. Syst. (TOCS)*, vol. 33, no. 3, pp. 7:1–7:55, Aug. 2015. [Online]. Available: <http://doi.acm.org/10.1145/2806887>
- [25] A. Sampson, J. Nelson, K. Strauss, and L. Ceze, "Approximate Storage in Solid-State Memories," *ACM Trans. Comput. Syst. (TOCS)*, vol. 32, no. 3, pp. 9:1–9:23, Sep. 2014. [Online]. Available: <http://doi.acm.org/10.1145/2644808>
- [26] S. Ahmadian, F. Taheri, M. Lotfi, M. Karimi, and H. Asadi, "Investigating Power Outage Effects on Reliability of Solid-State Drives," in *to appear in Design, Automation Test in Europe Conference Exhibition (DATE)*, March 2018.
- [27] Y. Cai, S. Ghose, E. F. Haratsch, Y. Luo, and O. Mutlu, "Errors in Flash-Memory-Based Solid-State Drives: Analysis, Mitigation, and Recovery," *arXiv preprint arXiv:1711.11427*, 2017.
- [28] Y. Cai, E. F. Haratsch, O. Mutlu, and K. Mai, "Error Patterns in MLC NAND Flash Memory: Measurement, Characterization, and Analysis," in *Proceedings of the Conference on Design, Automation and Test in Europe (DATE)*, 2012, pp. 521–526.
- [29] Y. Cai, S. Ghose, E. F. Haratsch, Y. Luo, and O. Mutlu, "Error Characterization, Mitigation, and Recovery in Flash Memory Based Solid-State Drives," *Proceedings of the IEEE*, vol. 105, pp. 1666–1704, 2017.
- [30] Y. Luo, S. Ghose, Y. Cai, E. F. Haratsch, and O. Mutlu, "Improving 3D NAND Flash Memory Lifetime by Tolerating Early Retention Loss and Process Variation," *Proceedings of the ACM on Measurement and Analysis of Computing Systems (POMACS)*, vol. 2, no. 3, p. 37, 2018.
- [31] D. A. Patterson, G. Gibson, and R. H. Katz, *A Case for Redundant Arrays of Inexpensive Disks (RAID)*. ACM, 1988, vol. 17, no. 3.
- [32] D. Arteaga, J. Cabrera, J. Xu, S. Sundararaman, D. Zhao, MingArteaga, J. Cabrera, J. Xu, S. Sundararaman, and M. Zhao, "Cloud-Cache: On-Demand Flash Cache Management for Cloud Computing," in *File and Storage Technologies (FAST)*, 2016.
- [33] M. Ryu, H. Kim, and U. Ramachandran, "Why are State-of-the-art Flash-Based Multi-Tiered Storage Systems Performing Poorly for HTTP Video Streaming?" in *Proceedings of the 22nd international workshop on Network and Operating System Support for Digital Audio and Video*. ACM, 2012, pp. 3–8.
- [34] B. Gregg. (2008) ZFS L2ARC. [Online]. Available: <http://blogs.sun.com/brendan/entry/test>
- [35] —. (2008) ZFS L2ARC. [Online]. Available: <http://www.brendangregg.com/blog/2008-07-22/zfs-l2arc.html>
- [36] Oracle. (2010) What Is ZFS? [Online]. Available: <https://docs.oracle.com/cd/E19253-01/819-5461/zfs-over-2/>
- [37] D. Jiang, Y. Che, J. Xiong, and X. Ma, "uCache: A Utility-Aware Multilevel SSD Cache Management Policy," in *High Performance Computing and Communications IEEE International Conference on Embedded and Ubiquitous Computing (HPCC_EUC)*, 2013.
- [38] Y. Liu, X. Ge, X. Huang, and D. H. Du, "MOLAR: A Cost-Efficient, High-Performance Hybrid Storage Cache," in *Cluster Computing (CLUSTER)*, *IEEE International Conference on*, 2013.
- [39] S. R. Dullloor, A. Roy, Z. Zhao, N. Sundaram, N. Satish, R. Sankaran, J. Jackson, and K. Schwan, "Data Tiering in Heterogeneous Memory Systems," in *Proceedings of the Eleventh European Conference on Computer Systems*. ACM, 2016, p. 15.
- [40] S. Lee, H. Bahn, and S. H. Noh, "CLOCK-DWF: A Write-History-Aware Page Replacement Algorithm for Hybrid PCM and DRAM Memory Architectures," *IEEE Transactions on Computers (TC)*, vol. 63, no. 9, pp. 2187–2200, 2014.
- [41] G. Dhiman, R. Ayoub, and T. Rosing, "PDRAM: A Hybrid PRAM and DRAM Main Memory System," in *Proceedings of the 46th Annual Design Automation Conference (DAC)*. ACM, 2009, pp. 664–469.
- [42] L. E. Ramos, E. Gorbatov, and R. Bianchini, "Page Placement in Hybrid Memory Systems," in *Proceedings of the International Conference on Supercomputing*. ACM, 2011, pp. 85–95.
- [43] Z. Zhang, L. Ju, and Z. Jia, "Unified DRAM and NVM Hybrid Buffer Cache Architecture For Reducing Journaling Overhead," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2016, pp. 942–947.
- [44] H. Liu, Y. Chen, X. Liao, H. Jin, B. He, L. Zheng, and R. Guo, "Hardware/Software Cooperative Caching For Hybrid DRAM/NVM Memory Architectures," in *Proceedings of the International Conference on Supercomputing*. ACM, 2017, p. 26.
- [45] S. Mittal and J. S. Vetter, "A Survey of Software Techniques For Using Non-Volatile Memories For Storage and Main Memory Systems," *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, vol. 27, no. 5, pp. 1537–1550, 2016.
- [46] X. Dong, N. Muralimohanar, N. Jouppi, R. Kaufmann, and Y. Xie, "Leveraging 3D PCRAM Technologies to Reduce Checkpoint Overhead For Future Exascale Systems," in *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*. ACM, 2009, p. 57.
- [47] J. C. Mogul, E. Argollo, M. Shah, and P. Faraboschi, "Operating System Support for NVM Hybrid Main Memory," 2009.
- [48] M. Gamell, I. Rodero, M. Parashar, and S. Poole, "Exploring Energy and Performance Behaviors of Data-Intensive Scientific Workflows on Systems With Deep Memory Hierarchies," in *International Conference on High Performance Computing (HiPC)*. IEEE, 2013, pp. 226–235.
- [49] H. A. Khouzani, C. Yang, and J. Hu, "Improving Performance and Lifetime of DRAM-PCM Hybrid Main Memory Through a Proactive Page Allocation Strategy," in *Asia and South Pacific Design Automation Conference (ASP-DAC)*. IEEE, 2015, pp. 508–513.
- [50] R. Rodriguez-Rodriguez, F. Castro, D. Chaver, L. Pinuel, and F. Tirado, "Reducing Writes in Phase-Change Memory Environments by Using Efficient Cache Replacement Policies," in *Proceedings of the Conference on Design, Automation and Test in Europe (DATE)*. EDA Consortium, 2013, pp. 93–96.
- [51] T. Hirofuchi and R. Takano, "RAMinate: Hypervisor-Based Virtualization for Hybrid Main Memory Systems," in *Proceedings of the Seventh ACM Symposium on Cloud Computing*. ACM, 2016, pp. 112–125.
- [52] S. Baek, H. G. Lee, C. Nicopoulos, and J. Kim, "Designing Hybrid DRAM/PCM Main Memory Systems Utilizing Dual-Phase Compression," *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, vol. 20, no. 1, p. 11, 2014.
- [53] L. Liu, H. Yang, Y. Li, M. Xie, L. Li, and C. Wu, "Memos: A Full Hierarchy Hybrid Memory Management Framework," in *IEEE International Conference on Computer Design (ICCD)*. IEEE, 2016, pp. 368–371.
- [54] SNIA. (2016) Microsoft Enterprise Traces, Storage Networking Industry Association IOTTA Repository. <http://h18000.www1.hp.com/products>. [Online]. Available: <http://h18000.www1.hp.com/products>
- [55] D. Narayanan, A. Donnelly, and A. Rowstron, "Write Off-Loading: Practical Power Management for Enterprise Storage," *ACM Transactions on Storage (TOS)*, vol. 4, no. 3, p. 10, 2008.
- [56] QEMU. (2016) QEMU: Quick Emulator. <http://wiki.qemu.org>. [Online]. Available: <http://wiki.qemu.org>
- [57] M. Wachs, M. Abd-El-Malek, E. Thereska, and G. R. Ganger, "Argon: Performance Insulation for Shared Storage Servers," in *Proceedings of the USENIX Conference on File and Storage Technologies (USENIX FAST)*, 2007.

- [58] C. A. Waldspurger, N. Park, A. Garthwaite, and I. Ahmad, "Efficient MRC Construction with SHARDS," in *File and Storage Technologies (FAST)*, 2015.
- [59] M. K. Qureshi, V. Srinivasan, and J. A. Rivers, "Scalable High Performance Main Memory System Using Phase-Change Memory Technology," *ACM SIGARCH Computer Architecture News*, vol. 37, no. 3, pp. 24–33, 2009.
- [60] J. Meza, J. Chang, H. Yoon, O. Mutlu, and P. Ranganathan, "Enabling Efficient and Scalable Hybrid Memories Using Fine-Granularity DRAM Cache Management," *IEEE Computer Architecture Letters*, vol. 11, no. 2, pp. 61–64, 2012.
- [61] Y. Li, S. Ghose, J. Choi, J. Sun, H. Wang, and O. Mutlu, "Utility-Based Hybrid Memory Management," in *IEEE International Conference on Cluster Computing (CLUSTER)*. IEEE, 2017, pp. 152–165.
- [62] N. Agarwal, D. Nellans, M. Stephenson, M. O'Connor, and S. W. Keckler, "Page Placement Strategies For GPUs Within Heterogeneous Memory Systems," *ACM SIGPLAN Notices*, vol. 50, no. 4, pp. 607–618, 2015.
- [63] N. Agarwal and T. F. Wenisch, "Thermostat: Application-Transparent Page Management For Two-Tiered Main Memory," *ACM SIGARCH Computer Architecture News*, vol. 45, no. 1, pp. 631–644, 2017.
- [64] B. Goglin, "Exposing the Locality of Heterogeneous Memory Architectures to HPC Applications," in *Proceedings of the Second International Symposium on Memory Systems*. ACM, 2016, pp. 30–39.
- [65] R. Salkhordeh and H. Asadi, "An Operating System Level Data Migration Scheme in Hybrid DRAM-NVM Memory Architecture," in *Proceedings of the Conference on Design, Automation & Test in Europe (DATE)*. EDA Consortium, 2016, pp. 936–941.
- [66] B. C. Lee, E. Ipek, O. Mutlu, and D. Burger, "Architecting Phase Change Memory As A Scalable DRAM Alternative," in *ACM SIGARCH Computer Architecture News*, vol. 37, no. 3. ACM, 2009, pp. 2–13.
- [67] B. C. Lee, P. Zhou, J. Yang, Y. Zhang, B. Zhao, E. Ipek, O. Mutlu, and D. Burger, "Phase-Change Technology and The Future of Main Memory," *IEEE Micro*, vol. 30, no. 1, 2010.
- [68] E. Kültürsay, M. Kandemir, A. Sivasubramaniam, and O. Mutlu, "Evaluating STT-RAM As An Energy-Efficient Main Memory Alternative," in *IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. IEEE, 2013, pp. 256–267.
- [69] N. Megiddo and D. S. Modha, "ARC: A Self-Tuning, Low Overhead Replacement Cache," in *File and Storage Technologies (FAST)*, vol. 3, no. 2003, 2003, pp. 115–130.
- [70] EnhanceIO. (2012) EnhanceIO. <https://github.com/stecinc/EnhanceIO>. [Online]. Available: <https://github.com/stecinc/EnhanceIO>
- [71] FIO. (2018) FIO: Flexible I/O Tester. <https://github.com/axboe/fio>. [Online]. Available: <https://github.com/axboe/fio>
- [72] R. McDougall. (2017) Filebench: Application Level File System Benchmark. <http://filebench.sourceforge.net>. [Online]. Available: <http://filebench.sourceforge.net>
- [73] Chen, Peter M and Lee, Edward K and Gibson, Garth A and Katz, Randy H and Patterson, David A, "RAID: High-Performance, Reliable Secondary Storage," *ACM Computing Surveys (CSUR)*, 1994.
- [74] Y. Chai, Z. Du, X. Qin, and D. A. Bader, "WEC: Improving Durability of SSD Cache Drives by Caching Write-Efficient Data," *IEEE Transactions on Computers*, vol. 64, no. 11, pp. 3304–3316, 2015.
- [75] R. L. Mattson, J. Gecsei, D. R. Slutz, and I. L. Traiger, "Evaluation Techniques for Storage Hierarchies," *IBM Systems journal*, 1970.
- [76] Q. Niu, J. Dinan, Q. Lu, and P. Sadayappan, "PARDA: A Fast Parallel Reuse Distance Analysis Algorithm," in *Parallel & Distributed Processing Symposium (IPDPS)*, 2012.
- [77] C. Ding and Y. Zhong, "Predicting Whole-Program Locality Through Reuse Distance Analysis," in *ACM SIGPLAN Notices*, 2003.
- [78] Y. Zhong, X. Shen, and C. Ding, "Program Locality Analysis Using Reuse Distance," *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 2009.
- [79] C. Fang, S. Carr, S. Önder, and Z. Wang, "Reuse-Distance-Based Miss-Rate Prediction on a Per Instruction Basis," in *Proceedings of the workshop on Memory system performance*, 2004.
- [80] C. Ding and Y. Zhong, "Reuse Distance Analysis," *University of Rochester, Rochester, NY*, 2001.
- [81] E. Berg and E. Hagersten, "StatCache: A Probabilistic Approach to Efficient and Accurate Data Locality Analysis," in *Performance Analysis of Systems and Software, IEEE International Symposium (ISPASS)*, 2004.
- [82] X. Shen, J. Shaw, B. Meeker, and C. Ding, "Locality Approximation Using Time," in *the Symposium on Principles of Programming Languages (POPL)*, 2007.
- [83] M. Zhao and D. Yeung, "Using Multicore Reuse Distance to Study Coherence Directories," *ACM Trans. Comput. Syst. (TOCS)*, vol. 35, no. 2, pp. 4:1–4:49, Jul. 2017. [Online]. Available: <http://doi.acm.org/10.1145/3092702>
- [84] M.-J. Wu and D. Yeung, "Efficient Reuse Distance Analysis of Multicore Scaling for Loop-Based Parallel Programs," *ACM Trans. Comput. Syst. (TOCS)*, vol. 31, no. 1, pp. 1:1–1:37, Feb. 2013. [Online]. Available: <http://doi.acm.org/10.1145/2427631.2427632>
- [85] M. Badamo, J. Casarona, M. Zhao, and D. Yeung, "Identifying Power-Efficient Multicore Cache Hierarchies via Reuse Distance Analysis," *ACM Trans. Comput. Syst. (TOCS)*, vol. 34, no. 1, pp. 3:1–3:30, Apr. 2016. [Online]. Available: <http://doi.acm.org/10.1145/2851503>
- [86] H. Roh, M. Shin, W. Jung, and S. Park, "Advanced Block Nested Loop Join for Extending SSD Lifetime," *IEEE Transactions on Knowledge and Data Engineering*, 2017.
- [87] Y. Liang, Y. Chai, N. Bao, H. Chen, and Y. Liu, "Elastic Queue: A Universal SSD Lifetime Extension Plug-in for Cache Replacement Algorithms," in *Proceedings of the 9th ACM International on Systems and Storage Conference*, 2016.
- [88] S. Huang, Q. Wei, D. Feng, J. Chen, and C. Chen, "Improving Flash-Based Disk Cache With Lazy Adaptive Replacement," *ACM Transactions on Storage (TOS)*, 2016.
- [89] J. Liu, Y. Chai, X. Qin, and Y. Xiao, "PLC-cache: Endurable SSD Cache For Deduplication-Based Primary Storage," in *Mass Storage Systems and Technologies (MSST)*, 2014.