Technical Report

Department of Computer Science
and Engineering
University of Minnesota
4-192 EECS Building
200 Union Street SE
Minneapolis, MN 55455-0159 USA

TR 10-026

Hot Data Identification for Flash Memory Using Multiple Bloom Filters

Dongchul Park and David Du

October 05, 2010

Hot Data Identification for Flash Memory Using Multiple Bloom Filters

Dongchul Park and David H.C. Du

Department of Computer Science and Engineering
University of Minnesota–Twin Cities, MN 55455, USA
park@cs.umn.edu, du@cs.umn.edu

Abstract

Hot data identification can be applied to a variety of fields. Particularly in flash memory, it has a critical impact on its performance (due to garbage collection) as well as its lifespan (due to wear leveling). Although this is an issue of paramount importance in flash memory, it is the least investigated one. Moreover, all existing schemes focus only or mainly on a frequency viewpoint. However, recency factor also must be considered as much importantly as the frequency for hot data identification. In this paper, we propose a novel hot data identification scheme adopting multiple bloom filters to efficiently capture finer-grained recency as well as frequency. In addition to this scheme, we propose a window-based direct address counting (named WDAC) algorithm to approximate an ideal hot data identification as our baseline. Unlike the existing baseline algorithm that cannot appropriately capture recency information due to its exponential batch decay, our WDAC algorithm using a sliding window concept can capture very fine-grained recency information. Our experimental evaluation with diverse realistic workloads including real SSD traces demonstrates that our proposed scheme outperforms the state-of-the-art hot data identification scheme. In particular, our scheme not only consumes less memory (50% less) and requires less computational overhead up to 58%, but also improves its performance up to 65%.

1. Introduction

A tendency toward flash-based Solid State Drives (SSD) now holds sway even in the enterprise servers as well as in personal computers, especially, laptops. This trend of the storage world results from the recent technological breakthroughs in flash memory and dramatic reduction of its price.

Flash memory is organized in units of blocks and pages. Each block consists of a fixed number (32 or 64) of pages. Reads and writes in flash memory are performed on a page basis, while erases operate on a block basis. Its most distinguishing feature is that it does not allow in-place updates (i.e., overwriting). The flash memory system cannot overwrite new data into existing data. Instead, the new data are written to clean spaces somewhere, and then the old data are invalidated for reclamation in the future (i.e., out-of-place updates).

In order to resolve this out-of-place update issue, Flash Translation Layer (FTL) has been developed and deployed to flash memory to emulate in-place update like block devices Gupta [2009], Lee [2007], Park [2008], Shin [2009], Wu [2006]. This layer enables users to utilize the flash memory like disks or main memory on top of existing conventional file systems without significant modification. Since FTL just makes an attempt to emulate flash memory like block devices, it cannot provide a fundamental solution to resolve this issue. Thus, as time goes on, the out-of-place update inevitably causes the coexistence of numerous invalid (i.e., outdated) and valid data. In order to reclaim the spaces occupied by the invalid data, a new recycling policy is required, which is a so-called garbage collection. However, the garbage collection can give rise to not only considerable amount of valid data copies to other clean spaces, but also data erases to reclaim the invalidated data spaces. Data erase (1,500μs) is the most expensive operation in flash memory compared to data read (25μs) and write (200μs) Agrawal [2008]. Consequently, a garbage collection results in a significant performance overhead as well as unpredictable operational latency. Flash memory exhibits another limitation: cells can be erased for a limited number of times (e.g., 10K-100K). Thus, frequent erase operations reduce the lifespan of the flash memory, which produces the wear leveling issue. The objective of wear leveling is to improve its lifetime by evenly distributing cell erases over the entire flash memory Agrawal [2008]. Those two critical issues are fundamentally based on the hot data identification.

We can simply classify the frequently accessed data as hot data. Otherwise, they are regarded as cold data. This definition is still vague and takes only *frequency* (i.e., the number of appearance) into account. However, there is another important factor–*recency* (i.e., closeness to the present)–to identify hot data. In general, many access patterns in workloads exhibit high temporal localities Chang [2005]; therefore, recently accessed data are more likely to be accessed again in near future. This is the rationale for the recency factor to be considered in hot data classification. Thus, if we redefine hot data while considering recency as well as frequency, it would be like this: if any recently accessed data will be frequently accessed again in near future, they will be regarded as hot data. Moreover, current hot data may turn into cold if they are no longer accessed frequently in the future and vice versa.

The definition of hot data, however, can be different for each application and also can be applied to a variety of fields. First, this can be primarily used in data caching Debnath [2009], Kim [2008]. By caching these hot data in the memory space in advance, we can significantly improve system performance. In addition to this buffer replacement in a cache, it is applied to B-tree indexing in sensor networks Nath [2007]. FlashDB is a flash-based B-tree index structure optimized for sensor networks. In FlashDB, the B-tree node can be stored either in read-optimized mode or in write-optimized mode, whose decision can be easily made on the basis of a hot data identification algorithm. Flash memory also adopts this classification algorithm particularly for a garbage collection and a wear leveling Boboila [2010], Chang [2007], Soundararajan [2010]. We can perform a garbage collection more efficiently by collecting and storing hot (i.e., frequently updated) data to the same block, which can reduce the garbage collection overhead. Moreover, we also improve flash lifespan by allocating hot data to the flash blocks with low erase count. In this paper, out of diverse applications of hot data classification, we focus on the garbage collection and wear leveling in the flash memory systems. Thus, we consider only write accesses, not read accesses, for the hot data identification.

Flash memory contains a fixed small amount of memory (SRAM) inside; so we need to exploit the SRAM efficiently. Since FTL needs the majority of this memory for a more efficient address translation, we have to minimize this memory usage for the hot data identification. This is one of the challenges to design an efficient hot data identification scheme. Moreover, computational overhead is another important issue since hot data identification has to be triggered whenever every write request is issued.

Although this hot and cold data identification is an issue of paramount importance in flash memory, it has been least investigated. Existing hot data identification scheme

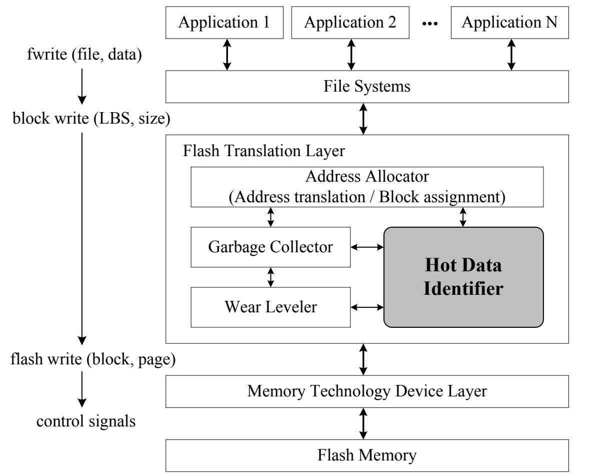

Figure 1. Typical System Architecture of Flash Memory-based Storage Systems

either suffers from large memory space requirements Chiang [1997] or incurs huge computational overhead to emulate LRU Chang [2002]. To overcome these problems, Hsieh et al. recently proposed a multiple hash function framework Hsieh [2006] to identify hot data. We refer to this scheme as multihash function scheme. This scheme adopts multiple hash functions and a counting bloom filter. Whenever a write request is issued, the LBA is hashed by multiple hash functions and then the corresponding counters in the bloom filter are increased by 1. If the counter value is greater than the predefined threshold (e.g., 4), the corresponding data are considered as hot. To implement an aging effect, it periodically (e.g., after every 5,117 write request) decreases all counter values by a half at a time (called exponential batch decay). This approach catches well frequency information because it maintains counters. However, it cannot appropriately capture recency information due to its batch decay.

While investigating other hot data identification schemes, we also posed a question about the existing baseline algorithm (i.e., direct address method in Hsieh [2006]): it cannot properly capture recency either. The direct address method assumes that unlimited memory space is available to keep track of hot data. It maintains a counter for each LBA to store access counts and periodically decays all LBA counting information thereby dividing by two at once. However, this algorithm also still has the same limitation as the multihash function approach since all LBAs accessed within its decay period are considered as having identical recency regardless of their access times or sequences.

Considering these observations, an efficient hot data identification scheme has to satisfy the following requirements: **1)** effective capture of recency information as well as frequency information, **2)** small memory consumption, and **3)** low runtime overhead. Based on these requirements, in this paper, we propose a novel hot data identification scheme based on multiple bloom filters. The key idea of this scheme is that each bloom filter has a different weight and recency coverage so that it can capture finer-grained recency information.

Whenever a write request is issued, the hash values of the LBA are recorded into one of multiple bloom filters (for short, BFs) in a round robin fashion. All information of each BF will be periodically erased by turns in order to maintain fine-grained recency information, which corresponds to an aging mechanism. Thus, each BF retains different a recency coverage. We also dynamically assign a different recency weight to each BF: the BF that just erased has higher recency weight, while the lowest recency weight is assigned to the BF that will be erased at right next turn because this BF has stored LBA access information for the longest period of time.

For frequency, our proposed scheme does not maintain a specific counter for all LBAs; instead, the number of BF recording the LBA information can exhibit its frequency information. The main contributions of this paper are as follows:

- ***An Efficient Hot Data Identification Scheme:*** A bloom filter can provide computational and space efficiency. Both multihash scheme and our proposed scheme try to take advantage of the bloom filter. Unlike the former using one counting bloom filter, the latter adopts multiple bloom filters. The multiple bloom filters enable our proposed scheme to capture finer-grained recency information so that we can achieve more accurate hot data classification. Multiple but smaller bloom filter empowers our scheme to require not only less memory space, but also lower computational overhead.
- ***A More Reasonable Baseline Algorithm:*** Our proposed approximation algorithm named Window-based Direct Address Counting (WDAC) adopts a window that is a conceptual buffer with a predefined size to store each coming request. Each LBA maintains a corresponding access counter. Whenever a write request is issued, the LBA is stored in the head of the window and the oldest one is evicted like a FIFO (First In First Out) queue. WDAC assigns all different recency weights to all LBAs in the window according to the closeness to the present. Thus, at the time a new request is arrived, all LBAs are shifted toward the tail of the window and all their recency values are reevaluated. Consequently, WDAC can catch precise (very fine-grained) recency as well as frequency.

The remainder of this paper is organized as follows. Section 2 gives an overview of flash memory and bloom filters. It also describes existing hot data identification schemes. Section 3 explains the design and operations of our proposed hot identification scheme and WDAC scheme. Section 4 provides a variety of our experimental results and analyses. Finally, Section 5 concludes the discussion.

2. Background and Related Work

In this section, we explain flash memory characteristics and introduce the existing hot data identification schemes with their pros and cons.

2.1 The Characteristics of Flash Memory

Figure 1 describes a typical system architecture of flash memory-based file systems. Both Memory Technology Device (MTD) and Flash Translation Layer (FTL) are two major parts of flash memory architecture. The MTD provides primitive flash operations such as read, write, and erase. The FTL plays a role in address translation between Logical Block Address (LBA) and its Physical Block Address (PBA) so that users can utilize the flash memory with existing conventional file systems without significant modification. A typical FTL is largely composed of an address allocator and a cleaner. The address allocator deals with address translation and the cleaner tries to collect blocks filled with invalid data pages to reclaim them for its near future use Chang [2004]. This garbage collection is performed on a block basis; so all valid pages in the victim block must be copied to other clean spaces before the victim block is to be erased.

Another crucial issue in flash memory is wear leveling. Wear leveling is to evenly distribute the number of erase count for each block since, unlike DRAM, each block cell in flash memory has a limited life span. That is, after a limited number of erases to each block, errors can start to occur. Currently, the most common allowable number of write per block is typically 10K for MLC (Multilevel-Cell) and 100K for SLC (Singlelevel-Cell) Bauer [1995]. Wear leveling is largely categorized into dynamic wear leveling and static wear leveling. Dynamic wear leveling tries to balance the block erase count by allocating hot data to the blocks with low erase counts. However, this will not be able to guarantee a complete balance of block erase counts since some blocks with cold data may not be erased forever. On the contrary, static wear leveling can evenly distribute the erase counts by migrating cold data into blocks with high erase counts. Dynamic wear leveling has a shorter life expectancy than static wear leveling since a complete balance of block erase counts is hard to achieve.

In summary, the motivation of wear leveling is to prevent any cold data from staying at any block for a long period of time. Its main goal is to minimize the variance among erase count values for each block so that the life span of flash memory is to be maximized.

2.2 Bloom Filters

The main goal of the bloom filter is to probabilistically test set membership with a space efficient data structure Bloom [1970]. That is, it is largely used to check if an element belongs to a member of a set. Since the space efficiency is an important factor for the bloom filer, the correctness

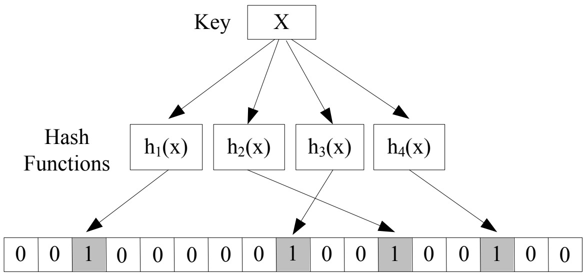

Figure 2. Bloom Filter Operations

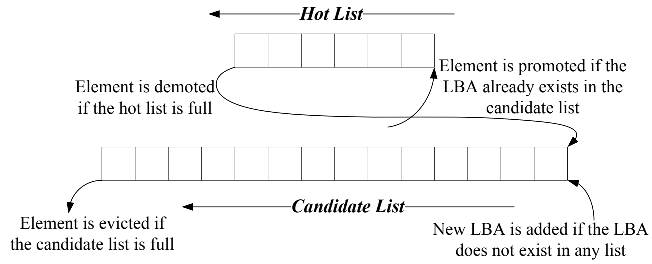

Figure 3. Two-level LRU List Scheme

can be sacrificed in order to maximize it. Thus, although a given key is not in the set, a bloom filter may provide a wrong positive answer called a *false positive*. However, the basic bloom filters never provide a false negative. We can also adjust design parameters (i.e., bloom filter size (M), the number of hash function (K) and the number of unique element (N)) of a bloom filter to allow a very low probability of the false positive.

The bloom filter is a bit array of M bits and all bits are initially set to 0. In addition to this bit array, there must also be K independent hash functions, each of which maps the given elements to the corresponding bit positions of the array thereby setting them all to 1. Figure 2 gives a specific example of a bloom filter with 4 independent hash functions. To insert an element to the bloom filter, the key value of the element is fed to the 4 hash functions so that we can get 4 hash values which correspond to the bit array positions. Then all the 4 bit positions are set to 1. Similarly all elements can be recorded to the bloom filter. To execute a membership query (i.e., check if the element is in the set), we first need to get the 4 bit positions by feeding the key value of the element to all 4 hash functions. If any of the 4 bits are 0, this means the element is not in the set because all the bits would have been set to 1 when the element was fed. If all are 1, there exist two possible cases: the corresponding element is in the set, or it is just a false positive due to the insertions of other elements. Assuming the probability of a false positive is very low, the answer to the query is positive.

The merit of a bloom filter is its low computational overhead since both insertions and membership test are executed in constant time. Moreover, we can achieve a good performance with high space efficiency, simple algorithm and efficient hash functions. However, it is difficult to store additional information with the keys in the bloom filter Bloom [1970].

2.3 Hot and Cold Data Identification Schemes

In this subsection, we describe the existing hot and cold data classification schemes and explore their advantages and disadvantages.

Chiang et al. Chiang [1997] proposed an out-of-update scheme called Flash Memory Server (FMS) in order to reduce the number of erase operations in flash memory. They made an attempt to classify data into three types for an efficient garbage collection: read-only, hot and cold data. Although FMS exhibits a good hot and cold data classification performance, it requires a large amount of memory space since it must keep the last access time information of all LBAs.

To resolve this limitation, Chang et al. Chang [2002] used a two-level LRU list that consists of a hot LBA list and a candidate list (Figure 3). Both lists are of fixed size and operate under LRU. Whenever a write request comes into the FTL driver, if the corresponding LBA already exists in the hot list, the data are classified as hot data. Otherwise, it is defined as cold data. If the LBA is not in the hot list but in the candidate list, the data are promoted to the hot list. If the data are not in either list, it will be inserted into the candidate list. This two-level LRU scheme consumes less memory than FMS; nevertheless, it incurs other problems. The performance of hot data identification is totally dependent on the sizes of both lists. In other words, a small hot list can save memory space, but its performance decreases because highly likely hot data may be demoted to the candidate list or even evicted from the candidate list. Moreover, this scheme requires high computing overheads to emulate LRU discipline.

Recently, Hsieh et al. Hsieh [2006] proposed a multiple hash function framework to identify hot data (Figure 4). This adopts multihash functions and one bloom filter with a D-bit counter for each bit position in the bloom filter to capture the frequency of the data access by incrementing the corresponding counters. The recency of the hot data is implemented by dividing the counter value by two periodically. If any one bit of B-most significant bits is set to 1, this bit position is considered as 1. Similarly if all K-bit positions (from the K-hash functions) are 1, the data are classified as hot. For example, as shown in Figure 4, assuming this scheme adopts 4-bit counters and 2-most significant bits, 4 will be its hot threshold value. Thus, the access counter values of corresponding positions are all greater than or equal to 4, the data in the LBA are identified as hot. After a specified time period, this scheme decreases all the counter values by a half with 1-bit right shifting to implement a decay effect. Compared to the other schemes, this achieves relatively less memory consumption as well as less computing overhead.

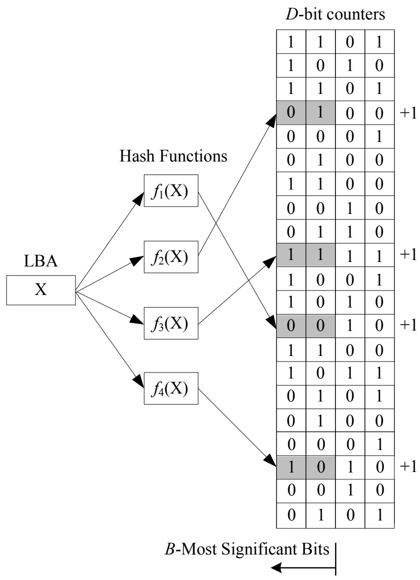

Figure 4. Multiple Hash Function Framework. Here, D = 4 and B = 2.

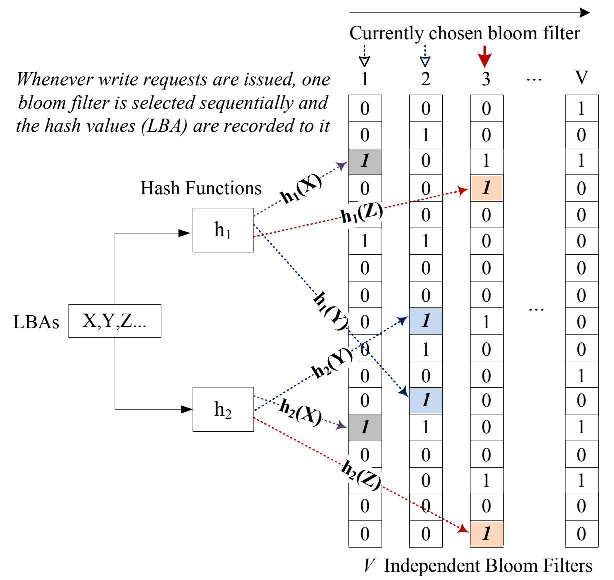

Figure 5. Our Framework and Its Operations

However, it does not appropriately catch recency information due to its exponential decrement of all LBA counters.

3. Multiple Bloom Filter-based Hot Data Identification Scheme

This section describes our proposed multiple bloom filter-based hot data identification scheme (Section 3.1) and our Window-based Direct Address Counting (WDAC) scheme as our baseline algorithm (Section 3.2).

3.1 The Framework

As shown in Figure 5, our scheme adopts a set of V independent bloom filters (for short, BFs) and K independent hash functions to capture both frequency and finer-grained recency. Each BF consists of M bits to record K hash values. The basic operation is simple: whenever a write request is issued to the Flash Translation Layer (FTL), the corresponding LBA is hashed by the K hash functions. The output values of each hash function ranges from 1 to M, and each hash value corresponds to a bit position of the M-bit BF respectively. Thus, K hash values set the corresponding K bits in the first BF to 1. When the next write request comes in, our scheme chooses the next BF in a round robin fashion to record its hash values. In addition, it periodically selects one BF in a round robin manner and erases all information in that BF to reflect a decay effect.

• ***Frequency:*** Unlike the multihash function framework adopting 4-bit counters, we do not maintain the specific BF counters for each LBA to count the number of appearance. Instead, our scheme investigates multiple BFs to check if each BF has recorded the corresponding LBA. The number of BF retaining the LBAs can show its frequency. However, since the frequency is closely associated with recency in our scheme, a different recency weight is assigned to each BF. This will be explained in the following subsection in more detail.

For precise frequency capturing, when it chooses one of the V BFs and marks the corresponding bits to 1, it does not allow overwriting to the BFs. That is, if the selected BF has already recorded the hash values of the LBA, it sequentially (in a round robin manner) examines other BFs until it finds a new one that has not recorded the LBA. This sequential examination minimizes disruption of our recency analysis. If it finds such a new one, it records the hash values to the BF. If it turns out that all (or the predefined number of) BFs have already contained the LBA information, our scheme simply defines the data as hot and skips its further processes such as a BF checking or a recency weight assignment since this will be definitely over the threshold. This shortcut decision reduces its overhead (this will be demonstrated in our experiment section). Consequently, assuming the hash values of an LBA appear in r $(0 \le r \le V)$ numbers of the BFs out of V BFs, we can say the corresponding LBA has appeared r times before. On the other hand, if all bloom filters have already recorded the LBA, we do not know its precise frequency number from then on. However, for hot data identification in our scheme, the information required is simply the fact if the value r is larger than the threshold, not the precise counter value. If this r is larger than a threshold, it passes the frequency check.

• ***Recency:*** In addition to the frequency, recency is another significant factor to make hot data decision. Even though the total access counters of two LBAs are equivalent in a period, the one accessed heavily in the further past and has never been accessed afterwards should be classified differ-

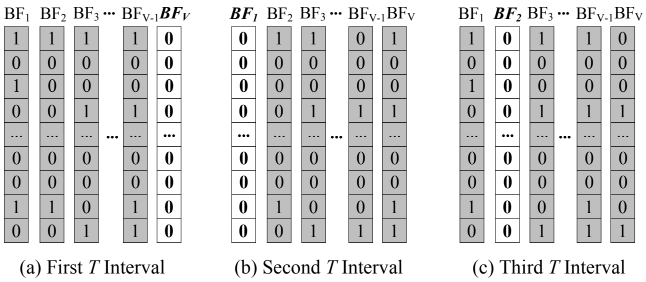

Figure 6. Our Aging Mechanism. Here, a white (i.e., not-shaded) bloom filter corresponds to a reset bloom filter

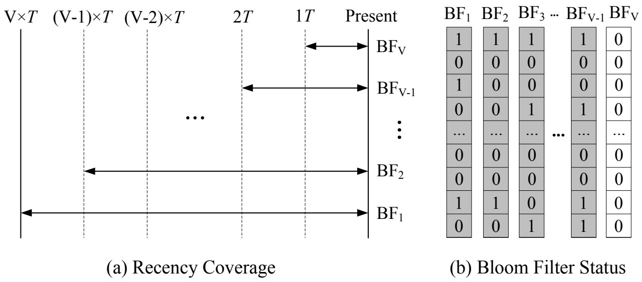

Figure 7. Recency Coverage for Each Bloom Filter

ently from the other one heavily accessed recently. If a hot data identifier depends only on counter values (frequency information), it cannot make distinction between them. Even though it has a decay mechanism, it cannot classify them well if the identifier captures coarse-grained recency.

Since our scheme does not maintain LBA counters, we devise a different aging mechanism to capture recency information. Figure 6 illustrates our aging mechanism to decay old information. Consider that we adopt *V* independent bloom filters and the hash values of LBAs are recorded to each bloom filter in a round-robin manner during a predefined interval *T*. An interval *T* represents a fixed number of consecutive write requests, not the time interval. As shown in Figure 6 (a), after the interval *T*, the BF that has not been selected in the longest time interval is selected and all bits in the BF (i.e., BF_V) are reset to 0 (i.e., all information of the bloom filter is erased). Note that since all BFs have served for the same period of time in the initial stage, we can randomly choose one of them and follow a round robin manner for the aging mechanism from then on. As soon as it is reset, the hashed LBA values start to be recorded again to all BFs including the reset BF. After the interval *T*, the next BF (BF_1) is selected and all the bits are reset (shown in Figure 6 (b)). Similarly, after the next *T* interval, the next BF (BF_2) is chosen in a right cyclic shift manner and all information is erased (shown in Figure 6 (c)) as time goes on.

Figure 7 shows the recency coverage after the interval *T* as soon as (BF_V) is reset. The reset BF (BF_V) can remember LBA information accessed during only the last one interval *T* (i.e., latest *1T* interval). The previously reset BF (BF_{V-1}) can record the LBA information accessed during the last two intervals. Similarly, the BF 1 (BF_1) which will be chosen as a next reset BF after this period can cover the longest interval $V \times T$. This means BF_1 records all LBA information for the last $V \times T$ intervals.

Our proposed scheme assigns a different recency weight to each BF so that recency value is combined with frequency value for hot data decision. The reset BF (BF_V) records most recent access information; so highest recency weight has to be assigned to it, whereas lowest recency weight is allotted to the BF that will be chosen as a next reset BF (BF_1) because it has recorded the information for the longest time. We intend to use recency value as a weight to the frequency value such that a final value combining both can be produced. Consequently, even though two different LBAs have appeared once in BF_V and BF_1 respectively, both frequency values are regarded differently as follows: assuming BF_V is a current reset BF, we first choose one BF ($BF_{\lfloor V/2 \rfloor}$) that has medium recency coverage and then assign a weight value of 1 to it. Next, we assign double weight value to the reset BF (BF_V), which means the LBA appearance in this BF counts as a double frequency (i.e., two appearances). Third, we calculate the difference of recency weights between each BF by $1/(V - \lfloor V/2 \rfloor)$ and assign evenly decreasing weight value by the difference from BF_{V-1} to BF_1. Finally each BF has different recency weights with evenly decreasing weight value order from BF_V to BF_1. For example, consider we adopt 4 BFs ($V = 4$) and BF_4 is a current reset BF, we first select BF_2 and assign a weight value of 1 to it. Then, we allot a double weight value (i.e., 2) to the reset BF (BF_4). Next, we can get 0.5 ($= 1/(4-2)$) as a difference of recency weights and finally distribute different recency weight values 2, 1.5, 1, and 0.5 to BF_4, BF_3, BF_2, and BF_1 respectively. Here, if any LBA appears both in BF_3 and BF_4, the original frequency value would be 2, but our scheme regards the final value as 3.5 because each BF has different recency weights. In this paper, we define this combination value of frequency and recency as a *hot data index*. Thus, if the value of hot data index is greater than or equal to a predefined threshold value, we regard the data as hot. Algorithm 1 describes a regular checkup operation in our multiple BF-based scheme.

• ***A Bloom Filter Size and Decay Period:*** Applications of a BF are required to configure three parameters: a BF size (*M*), the number of hash function (*K*) and the number of unique element (*N*). These three parameters are closely related with each other and have a crucial impact on the performance of its applications.

According to Dharmapurikar [2006], we can estimate an optimal BF size (*M*) using the following formula, $M = KN/ln2$, given *K* and *N*. It is clear that the BF size (*M*) needs to be quite large compared to the size of the element set (*N*). For example, assuming we adopt two hash functions ($K = 2$) and there are 4,096 unique elements, we

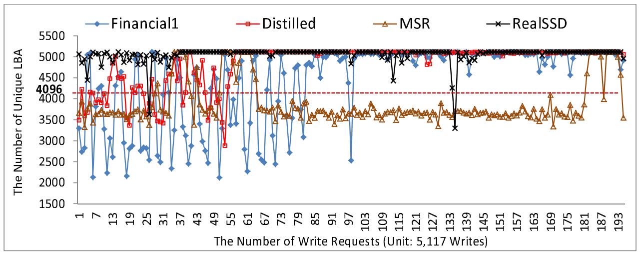

(a) The Number of Unique LBAs within Unit Period (5,117 Writes)

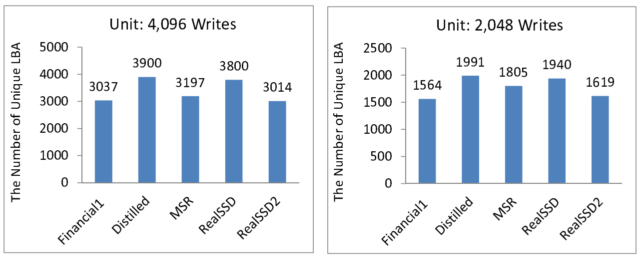

(b) Average No. of Unique LBAs within Each Unit Period under Various Workloads

Figure 8. The Number of Unique LBAs within Unit Periods under Various Workloads

should adopt an 11,819-bit vector as an appropriate size of the BF to minimize a false positive probability. This basic BF allows the elements to be only added to it, but not removed from it. Thus, as the number of input elements grows, the data recorded in the BF are also continuously accumulated, which causes a higher false positive probability. To reduce this, the traditional BF scheme requires a large BF size as described in the formula. However, we cannot simply apply this relationship to both our proposed scheme and multihash function scheme. The BFs in both schemes retain a distinguishing feature from the basic BF–information in the bloom filters can be removed. Especially our proposed scheme completely removes all information in one of the V BFs periodically. This prevents continuous data accumulation to the BF. Similarly, the multihash scheme periodically decays some of the information in the filter by one-bit right shift manner. Consequently, we can adopt a smaller BF than the aforementioned traditional BF.

Multihash function scheme consists of 4,096 entries of a BF each of which is composed of a 4-bit counter. It also adopts 5,117 write requests (N) as its decay period. This is based on their expectation that the number of hash table entry (M) can accommodate all those LBAs which corresponds to cold data within every N (where, $N \leq M/(1-R)$) write requests (here, R is the hot ratio of a workload and the authors assumed R is 20%) Hsieh [2006]. To verify their assumption, as displayed in Figure 8 (a), we measured the number of unique LBA for every 5,117 write request under several real traces such as Financial1, Distilled, MSR and RealSSD (these traces will be explained in our experiment section). Figure 8 (a) clearly shows that the number of unique LBA is very frequently over than their BF size (4,096) for each unit period (every 5,117 write request) under all four traces. This necessarily causes hash collisions so that it results in a higher false positive probability.

Algorithm 1 A Regular Checkup Operation

Function HotColdCheckup(BloomFilter)

```
V = The Number of Bloom Filter
K = The Number of Hash Function
RECENCY_WEIGHT = 1 / (V − ⌊V/2⌋)
Write Request is issued with LBA X
RequestCounter++
// LBA X is fed to each hash function
for (i=1; i≤K; i++) do
    l_i = h_i(X)
end for
Choose one bloom filter, CurrentBF
if (All K-bit positions in CurrentBF are already set to 1) then
    Try to choose another bloom filter sequentially
    if (Found one available) then
        for (i=1; i≤K; i++) do
            CurrentBF[l_i] = 1
        end for
    end if
    // All bloom filters have already recorded the LBA
    if (Cannot find any bloom filter available) then
        CurrentBF = The Next Bloom Filter
        // Shortcut decision
        Define this as HOT data and return
    end if
else
    for (i=1; i≤K; i++) do
        CurrentBF[l_i] = 1
    end for
end if
CurrentBF = The Next Bloom Filter
for (i=1; i≤V; i++) do
    Check which bloom filters have already recorded the LBA before
    TotalHotDataIndex += CalculateTotalHotDataIndex(Hot_BF_Index)
end for
if (TotalHotDataIndex ≥THRESHOLD) then
    Define this as HOT data
end if
// Decay process
if ((RequestCounter % RESET_PERIOD)) then
    Reset all bits in reset bloom filter, ResetBF
    ResetBF = The Next Bloom Filter
end if
Function CalculateTotalHotDataIndex(Index)
Diff = The Index of ResetBF – The Index of Hot Bloom Filter
if (Diff ≥ 0) then
    HotDataIndexValue += (2 – (RECENCY_WEIGHT × Diff))
else
    HotDataIndexValue += (RECENCY_WEIGHT × |Diff|)
end if
Return HotDataIndexValue
```

We also measured the average number of unique LBA within both 2,048 and 4,096 write request periods under the same workloads. As shown in Figure 8 (b), all the numbers must be less than the number of each unit write request since they are closely related with average hot ratios of each workload: intuitively the higher average hot ratio, the less number of unique LBA. Based on these observations, as our decay interval T, we adopt M/V numbers of write requests, where M is a BF size and V is the number of BF. To reduce the probability of a false positive, the BF size (i.e., the number of hash table entries) must be able to accommodate at least all unique LBAs that came in for the last $V \times T$ interval. Thus, whenever M/V numbers of write requests are issued, one of the BFs are selected in a round robin fashion and all the bits in the BF are reset to 0. For example, assuming our BF size is 2,048, our decay

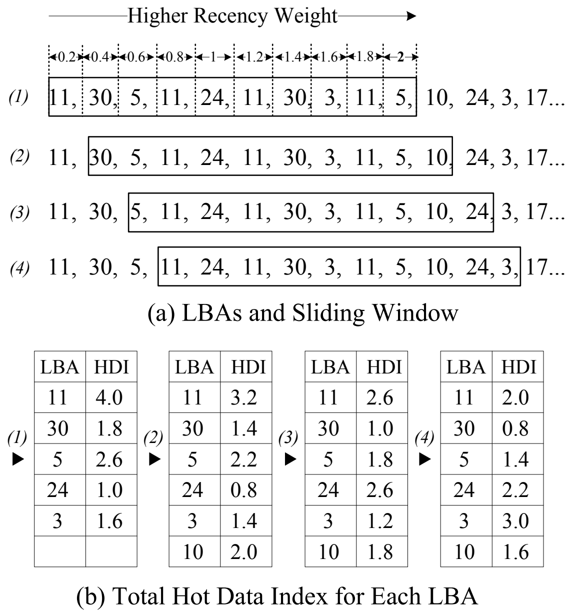

Figure 9. Working Process of WDAC Algorithm. Here, the window size is 10. HDI corresponds to the total hot data index value.

period corresponds to 512 write requests. Since memory consumption is also very important factor to design hot data identification scheme, we will adopt even less (a half) size of BF and show much better performance than the multihash function scheme.

3.2 WDAC: A Window-based Direct Address Counting

Ideally, if we know the future accesses of data, we will be able to easily identify the data as hot or cold data at present. However, we never know the future accesses in advance, which is the rationale of an approximation to the ideal hot data identification algorithm based on the past accesses.

Hsieh et al. proposed an approximated hot data identification algorithm named a direct address method (hereafter, we refer to this as DAM) as their baseline. We, however, raise a question regarding this. The DAM maintains LBA access counters. Whenever write requests are issued, the corresponding LBA counters are increased by 1 accordingly. By using counters it can capture the frequency information well, whereas, with respect to recency, it periodically forces all access counters to be reduced by a half at once. This aging mechanism, like the multihash function scheme, cannot reflect appropriate recency to each accessed LBA because all LBAs accessed within the same decay period have an identical recency regardless of their access time or sequence. In other words, it assigns the identical recency weight to all LBAs only if they are accessed within the same decay period. For instance, assuming the decay period corresponds to 4,096 write requests, the LBA accessed in the first request within this period is considered as having the same recency as the LBA accessed in the last $(4,096^{th})$ request.

To resolve this limitation, in this subsection, we propose a more reasonable baseline algorithm to approximate ideal hot data identification named Window-based Direct Address Counting (WDAC). As shown in Figure 9 (a), WDAC maintains a specific size of buffer like a sliding window. In addition, it maintains total hot data index values for each LBA (shown in Figure 9 (b)).

Within this window, all elements have a different recency value according to their access sequences: the closer to the present, the higher recency weight is assigned to the LBA. That is, highest recency value (i.e., 2) is assigned to the most recently accessed element (head) in the window, whereas, the lowest recency is allotted to the last element (tail) in the window. All intermediate LBAs in the window have all different recency values with an evenly decreasing manner. Assuming the window size is *W*, we can get the recency difference between two adjacent elements in the window as $2/W$. Thus, all recency weights assigned to each LBA evenly decrease by $2/W$ from the head to the tail in the window. Whenever a new LBA comes in, all recency values are reassigned to all the LBAs shifted in the window and the last one is evicted from the window.

As our highest recency weight, we choose 2. However, the average of all recency values in the window is equivalent to that within the same decay period in the DAM. Instead, we assign a higher recency weight to the recently accessed LBAs and vice versa. Consequently, our proposed WDAC can properly identify hot data thereby using very fine-grained recency information.

Figure 9 illustrates a simple example for our WDAC process. This window can contain 10 LBAs and recency difference corresponds to 0.2 ($= 2/10$). Whenever a write request comes in to the flash translation layer (FTL), this LBA is stored in the head of the window and the last one is evicted. At the same time, all the others are shifted toward the tail of the window and their recency values are reevaluated by decreasing their values by the recency difference (0.2). All of the total hot data index values also need to be updated accordingly. Finally, if the total hot data index value of the corresponding LBA is greater than or equal to a predefined hot threshold, we regard it as hot data.

In our proposed WDAC scheme, since window size is associated with recency and frequency, a proper window size is an important parameter to this scheme. Therefore, the impact of this window size will be discussed in our experiment part.

4. Experimental Results

This section provides diverse experimental results and comparative analyses. Furthermore, we explore and discuss the impact of both memory size, window size, and the number of a bloom filter for a more complete evaluation.

Table 1. System Parameters and Values

System Parameters	MBF	MHF	WDAC	DAM
Bloom Filter Size	2^{11}	2^{12}	N/A	N/A
Number of Bloom Filter	4	1	N/A	N/A
Decay (Window Size)	2^{9}	2^{12}	2^{12}	2^{12}
Number of Hash Function	2	2	N/A	N/A
Hot Threshold	4	4	4	4
Recency Weight Difference	0.5	N/A	0.000488	N/A

Table 2. Workload Characteristics

Workloads	Total Requests	Request Ratio (Read:Write)	Inter-arrival Time (Avg.)
Financial1	5,334,987	R:1,235,633(22%) W:4,099,354(78%)	8.19 ms
MSR	1,048,577	R:47,380(4.5%) W:1,001,197(95.5%)	N/A
Distilled	3,142,935	R:1,633,429(52%) W:1,509,506(48%)	32 ms
RealSSD	2,138,396	R:1,083,495(51%) W:1,054,901(49%)	492.25 ms

4.1 Evaluation Setup

We compare our Multiple Bloom Filter-based scheme (hereafter, refer to as MBF) with three other hot data identification schemes: Multiple Hash Function scheme (hereafter, refer to as MHF) Hsieh [2006], Direct Address Method (refer to as DAM) Hsieh [2006], and our proposed baseline scheme named WDAC. We, however, focus particularly on the MHF since it is the state-of-the-art scheme. We adopt a freezing approach as a solution for a counter overflow problem in MHF since it, in our experiments, showed a better performance than the other approach (i.e., exponential batch decay). DAM is an aforementioned baseline proposed in MHF scheme. Table 1 shows system parameters and their values. For fair evaluation, we assign the same number of write requests (4,096) for a decay interval in DAM as well as for a window size in our proposed WDAC algorithm. Moreover, we employ 4,096 write requests for the decay interval in MHF scheme to synchronize with DAM and WDAC, while we adopt 512 write requests for our decay period (this is described in Subsection 3.1). We also adopt identical hash functions for both MBF and MHF schemes. Lastly, since window size is 4,096 in WDAC, there exist about 0.488×10^{-3} ($= 2/4,096$) weight difference for each element and we also assign 0.5 weight difference for each BF in MBF.

For more objective evaluation, we adopt four real trace data (Table 2). Financial1 is a write intensive trace file from the University of Massachusetts at Amherst Storage Repository UMass [2002]. This trace file was collected from an Online Transaction Processing (OLTP) application running at a financial institution. Distilled trace file Chang [2002] shows a general and personal usage patterns in a laptop such as web surfing, watching movies, playing games, documentation work, etc. This is from the flash memory research group repository at National Taiwan University, and since the MHF scheme uses only this trace file, we also adopt this trace for fair evaluation. We also adopt MSR trace file made up of 1-week block I/O traces of enterprise servers at Microsoft Research Cambridge Lab Microsoft [2007]. We choose, in particular, *prxy volume 0* traces since it exhibits a write intensive workload Narayanan [2008]. Lastly, we employ a real Solid State Drive (SSD) trace file that is 1-month block I/O traces of a desktop computer (AMD X2 3800+, 2G RAM, Windows XP Pro) in our lab (hereafter, refer to as RealSSD trace file). We installed Micron's C200 SSD (30G, SATA) to the computer and collected personal traces such as computer programming, running simulations, documentation work, web surfing, etc.

The total requests in Table 2 correspond to the total number of read and write requests in each trace. These requests can also be subdivided into several or more sub-requests with respect to LBA accessed. For example, let us consider such a write request as *WRITE 100, 5*. This means 'write data into 5 consecutive LBAs from the LBA 100'. In this case, we regard this request as 5 write requests in our experiments.

4.2 Performance Metrics

A hot ratio is a ratio of hot data to all data. First of all, we choose this *hot ratio* to compare each performance of those four hot data identification schemes and to examine closeness among them. However, even though both hot ratios of two algorithms are identical, hot data classification results of both schemes may be able to be different since an identical hot ratio means the same number of hot data to all data and does not necessarily mean all classification results are identical. Thus, in order to make up for this limitation, we employ another evaluation metric: *false identification rate*. Whenever write requests are issued, we try to compare each identification result of each scheme. This enables us to make a more precise analysis of them. *Memory consumption* is another important factor to be discussed since SRAM size is very limited in flash memory. Finally, *runtime overhead* also must be taken in account. To evaluate it, we choose two main operations and measure CPU clock cycles per operation in each scheme.

4.3 Results and Analysis

We discuss our evaluation results in diverse respects.

• ***Baseline Algorithm:*** We first start to evaluate two baseline schemes: our proposed WDAC and DAM. As shown in Figure 10 (a) and (b), the hot ratios of both schemes exhibit considerably different results under Financial1 and MSR, whereas they display almost identical patterns under Distilled and RealSSD traces. However, as mentioned before, the identical hot ratios do not necessarily mean the same performance of them. Thus we make another experiment for a more comparative analysis.

Figure 11 shows the number of different identification results between WDAC and DAM. We count the number

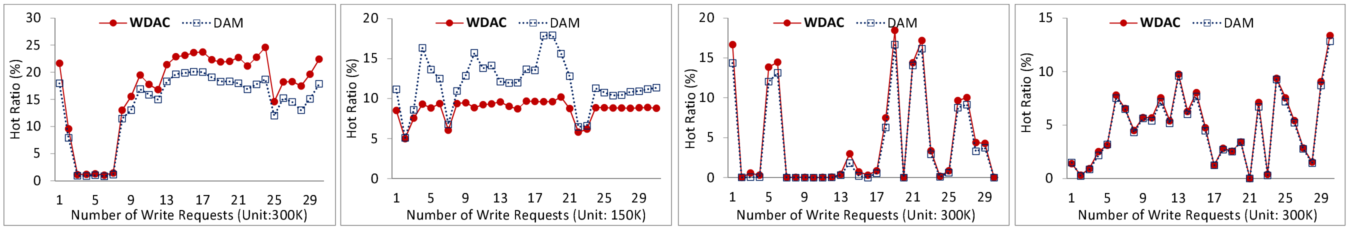

(a) Financial1 (b) MSR (c) Distilled (d) RealSSD

Figure 10. Hot Ratios of Two Baseline Algorithms

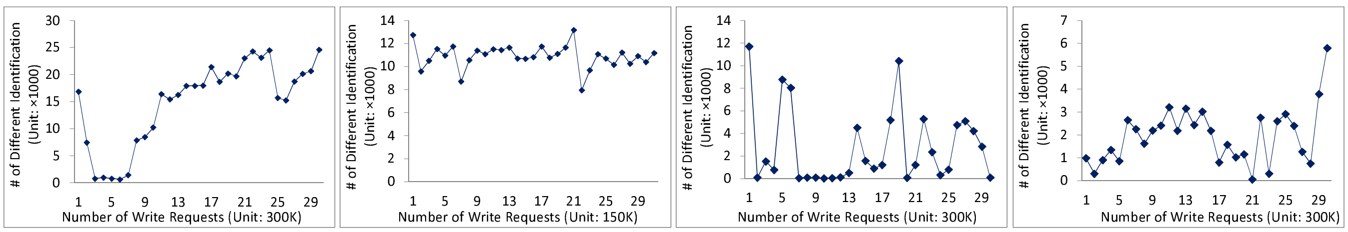

(a) Financial1 (b) MSR (c) Distilled (d) RealSSD

Figure 11. The Number of Different Identification between Two Baseline Algorithms

of different hot classification results during a specified unit time (150K or 300K write requests) between them. Thus, zero value means all the identification results are identical between each scheme during the corresponding unit time. Particularly, Figure 11 (c) and (d) illustrate well the performance of both schemes can be different even though two hot ratios look very similar each other as shown in Figure 10 (c) and (d).

• ***Our Scheme(MBF) vs. MHF:*** Now, we make an attempt to evaluate performance of our proposed scheme (MBF) and the multiple hash function scheme (MHF). For reference, we include our aforementioned two baselines.

As illustrated in Figure 12, our MBF presents a very close approach to our baseline, even to DAM, while the MHF has a tendency to exhibit considerably higher hot ratios than MBF as well as two baselines over all traces. This results from a higher ratio of false identification in MHF. That is, since the BF size (i.e., hash table size) is limited, even though write requests of non-hot data are increased, the chance of incorrect counter increments also grows due to hash collisions. As a result, it causes higher hot ratios. As Hsieh et al. are mentioned in their paper, the MHF scheme does not show a good performance (i.e., the false identification grows significantly) especially when the hot data ratio in traces is very low. To get over this limitation, they suggested two solutions: a larger size of the hash table or a more frequent decay operation. However, both suggestions cannot be fundamental solutions since not only does a larger hash table require more memory consumption, but also a more frequent decay produces a significant performance overhead.

Our scheme, on the other side, resolves that limitation thereby adopting multiple BFs. Unlike MFH adopting the

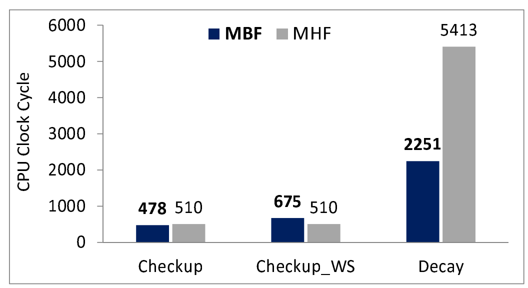

Figure 14. Runtime Overhead per Operations

exponential batch decay after a longer decay period (here, 4,096 write requests), our scheme erases all information only in one BF out of V BFs (here, V=4) after a shorter decay period (here, 512 write requests). Furthermore, a smaller BF results in lower computational overhead as well as less memory space consumption. Both memory consumption and runtime overhead will be discussed in the next subsections in more detail.

Figure 13 shows the false identification rates of both our scheme and MHF scheme. We compare both MBF and MHF with our proposed ideal scheme (WDAC). Thus, we can call these results false identification rates. We count the number of different hot classification results between two schemes and get a ratio of false identification results to a specified unit request period. As presented in Figure 13, our MBF scheme exhibits much lower false identification rates than MHF. It improves its performance by an average of 41%, 65%, 59%, and 36% under the four workloads respectively.

• ***Runtime Overhead:*** Computational overhead is another important factor to evaluate hot data identification scheme. Figure 14 displays runtime overheads for a checkup and de-

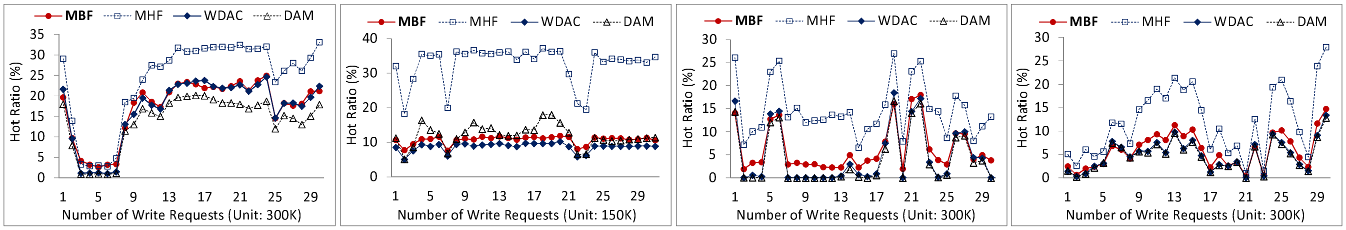

(a) Financial1 (b) MSR (c) Distilled (d) RealSSD

Figure 12. Hot Ratios of Four Schemes under Various Traces

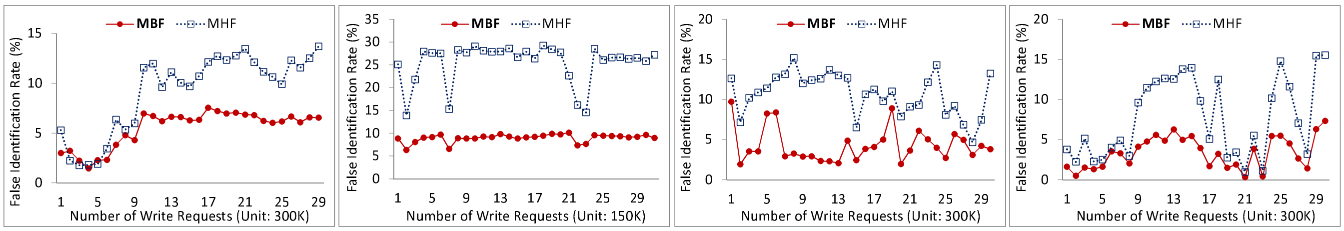

(a) Financial1 (b) MSR (c) Distilled (d) RealSSD

Figure 13. False Identification Rates of Both MBF and MHF

cay operation in MBF and MHF. A checkup operation means the verification process to check if the data in the corresponding LBA are hot whenever write requests are issued. A decay operation is the aforementioned aging mechanism. Since both are the most representative operations in hot data identification schemes, we choose and evaluate them. To evaluate each operation, we measure CPU clock cycles for them under the configurations in Table 1. This measuring is done over an AMD X2 3800+ (2GHz) system with 2G RAM under Windows XP Professional platform.

For fair and precise evaluation of each operation in both schemes, we feed many write requests from MSR traces (100K numbers of write requests) and then measure their average CPU clock cycles since a CPU clock cycle is dependent on cache misses (such as code miss and data miss) in CPU caches. In other words, when we measure CPU clock cycles of the operations, a lot more clock cycles are required at the beginning stage due to CPU cache misses, while the clock cycles start to be significantly reduced and stabilized soon afterwards.

As presented in Figure 14, the runtime overhead of a checkup operation in MBF is comparable to the checkup operation in MHF because both schemes adopt identical hash functions and BF data structures. However, our scheme requires about 6% less computational overhead than MHF since our scheme (MBF) sets only one bit for each hash function, while MHF frequently needs to set more than one bit. That is, since MHF scheme consists of 4-bit counters, one or more bits can need to be set in order to increase the corresponding counters. In the case that a chosen BF in MBF has already recorded the corresponding LBA information, it tries to choose another BF until it finds one available. Intuitively, this process may require extra overhead in comparison with the MHF. However, when it turns out that all BFs (here, 4 BFs) have already included the LBA information, MBF simply defines the data as hot and skips further processes such as a BF checking or recency weight assignment. This shortcut decision in the checkup operation considerably reduces its overhead by an average of 29% compared to the checkup operation without shortcut decision (478 *vs.* 675) in our scheme.

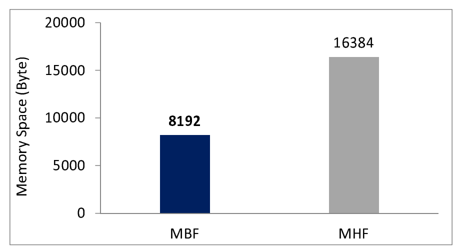

Figure 17. Memory Space Requirements

On the other hand, a decay operation of our scheme outperforms that of MHF. Our scheme consists of 4 numbers of 2,048 sizes of 1-bit array, while MHF is composed of 4,096 sizes of 4-bit array. Thus, whenever a decay operation is executed in MHF, all 4,096 numbers of 4-bit counters must be right-shifted by 1-bit (i.e., divided by 2) at once. Our scheme, however, needs to reset only 2,048 sizes of 1-bit array. Although our scheme requires more frequent (here, 4 times more frequent) decay operations than MHF, it requires almost a half (58%) less runtime overhead than MHF.

• ***Impact of Memory Size:*** Since SRAM is very limited in flash memory and expensive (generally 5$\sim$10 times more expensive than DRAM), we must consider memory consump-

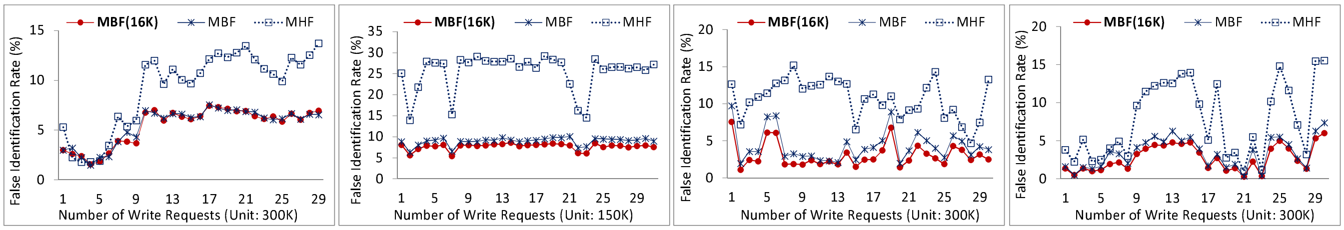

(a) Financial1 (b) MSR (c) Distilled (d) RealSSD

Figure 15. False Identification Rates between Two Schemes with Same Memory Space. Here, original MBF and MHF requires 8KB and 16KB respectively.

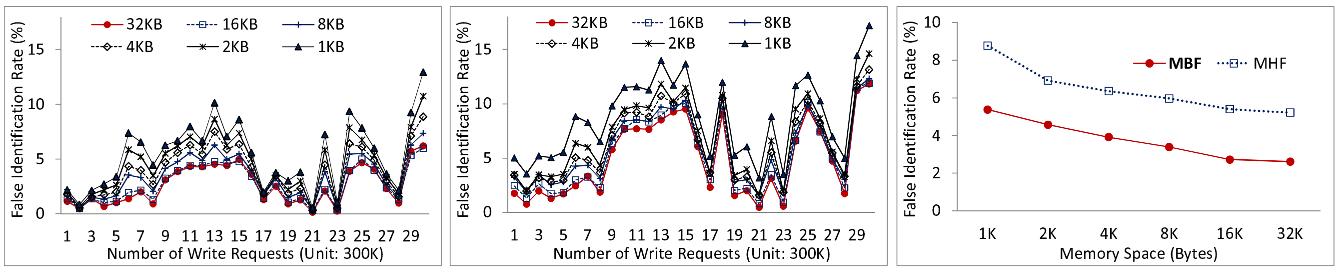

(a) Our Scheme (MBF) (b) MHF (c) Change of Total False Identificatio Rates

Figure 16. Performance Change over Various Memory Spaces under RealSSD Traces

tion to design a hot data identification scheme. Thus, evaluating performance of both MBF and MHF with the same size of memory space is also valuable. In addition, we explore the performance changes of both schemes over various memory sizes in this subsection.

Figure 17 displays memory space consumption of both schemes. MBF consumes only 8KB, while MHF requires a double (16KB). As mentioned before, our proposed scheme comprises 4 BFs each of which consists of 2KB (i.e., 2,048$\times$1-bit). On the other hand, MHF is composed of 4K numbers of 4-bit counters (i.e., 4KB$\times$4-bit). Therefore, our scheme consumes only a half of memory space of MHF scheme.

Figure 15 illustrates false identification rates between two schemes with same memory space. Since MHF originally requires 16KB, we double each BF size in our scheme from 2KB to 4KB (i.e., 4KB$\times$4=16KB). For reference, we include our original MBF (8KB) into each plot. As shown in Figure 15, MBF with 16KB improves its performance by lowering its false identification rates further than original MBF (8KB), which means our proposed scheme clearly outperforms MHF under the same memory space condition. MBF with 16K benefits from a larger BF size so that it can reduce the possibility of a false positive in each BF.

Different memory spaces (i.e., different bloom filter sizes) will have an effect on the performance of each scheme. We now explore the impact of memory space on both MBF and MHF. Figure 16 (a) and (b) exhibit performance changes of both schemes over various memory spaces under RealSSD trace file. Both schemes with larger memory space show a better performance (i.e., lower false identification rate) than those with smaller memory space. Figure 16 (c) shows more clearly the impact of memory space. We attempt to evaluate total false identification rates for each scheme under the different memory spaces. As a memory space grows, although each performance of both schemes is also improved accordingly, our proposed scheme still exhibits better performance than MHF throughout all memory spaces from 1KB to 32KB. Since other experimental results under the other traces such as Financial1, MSR and Distilled also show almost identical patterns with this experiment, we do not provide the other ones.

• ***Impact of Window Size:*** In WDAC scheme, since window size is closely related with hot data identification results, investigating the impact of window size is also valuable.

Intuitively, a larger window will cause a higher hot ratio because a window size is directly associated with frequency information. All experiments in Figure 18 demonstrate well this intuition. Since total hot data index values for each LBA are the summation of all weighted frequency values of the corresponding LBA within the window, the frequency value has a great impact on the hot data decision. A larger window can contain a more number of LBA access information so that it necessarily causes a higher frequency value for each LBA. This results in a higher hot ratio. As displayed in Figure 18, WDAC with a larger window size exhibits a higher hot ratio.

Conceptually the decay period in DAM is similar to the window size in WDAC. Thus, we also explore the relationship between various decay periods and hot ratios in DAM.

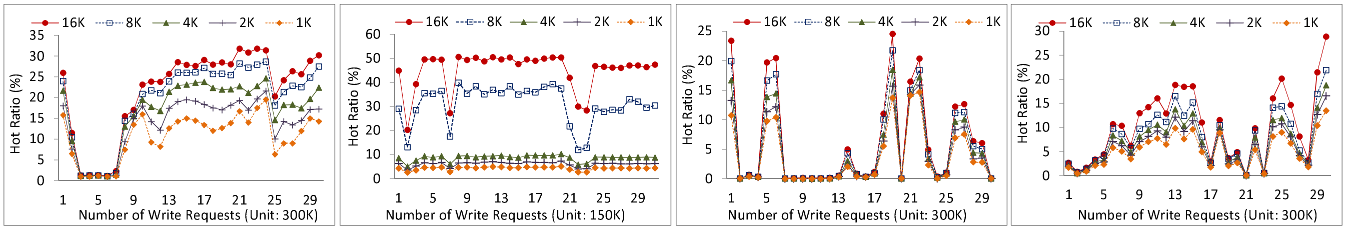

(a) Financial1 (b) MSR (c) Distilled (d) RealSSD

Figure 18. Changes of Hot Ratios over a Different Window Size in WDAC

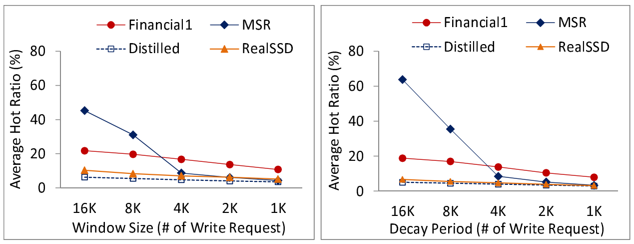

(a) WDAC (b) DAM

Figure 19. Changes of Average Hot Ratios over Various Window Sizes in WDAC and Decay Periods in DAM

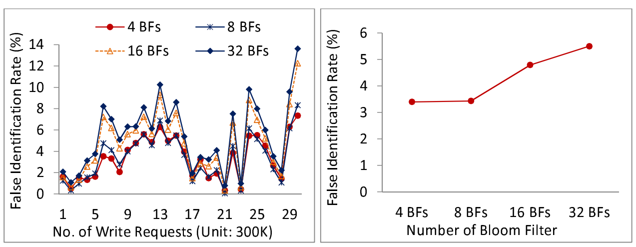

(a) MHF (b) Change of Total False Identification Rates

Figure 20. Performance Change over Various Numbers of Bloom Filter in Our Scheme under RealSSD trace

Figure 19 presents the changes of average hot ratios over a different window size in WDAC and a different decay period in DAM. Like WDAC, a longer decay period causes a higher hot ratio in DAM. Thus, overall trends of hot ratios in DAM exhibit very similar results.

• ***Impact of the Number of a Bloom Filter:*** In our proposed scheme, the number of BF corresponds to the granularity of recency. To explore its impact, we make experiments with various numbers of BF in our scheme. For more objective comparison, we not only assign the same memory space (8KB) and the same number of hash function (2) to each scheme, but also configure BF sizes and decay periods accordingly. As presented in Figure 20 (b), as the number of BF grows, the false identification rates also increase. This result is closely related with BF sizes of each scheme. Since each configuration has the same total memory consumption, smaller BFs have to be assigned to the scheme with a more number of BF, which results in higher false identification rates. Even though the scheme with a more number of BF can capture finer-grained recency information, this benefit is offset by its higher false identification rates. To verify this, we make another experiment with the same configurations except for memory consumption. When we double the number of BF from 4 to 8 while remaining the same BF size (i.e., 8KB to 16KB), the performance is improved by an average of 18% since it can benefit from capturing its finer-grained recency.

5. Conclusion

In this paper, we proposed a novel hot data identification scheme for flash memory-based storage systems. Unlike the multihash framework, our scheme adopts multiple bloom filters and each bloom filter has a different weight and a different recency coverage so that it can capture finer-grained recency information. Furthermore, multiple but smaller bloom filters empower our scheme to achieve not only lower runtime overheads, but also less memory consumption.

In addition to this novel scheme, we proposed a more reasonable baseline algorithm to approximate an ideal hot data identification named Window-based Direct Address Counting (WDAC). The existing algorithm (i.e., direct address method) cannot properly catch recency information because it assigns an identical recency weight to all LBAs accessed within a decay period. However, our WDAC allots all different recency weights to all LBAs within a window according to their access sequences so that it can capture precise recency as well as frequency information. Consequently, WDAC can properly identify hot data.

We made experiments in many respects under diverse real traces including real SSD traces. All hot data identification results (i.e., hot ratios) of our scheme display much closer results to those of the baseline scheme than the other one. Furthermore, to make up for the limitation of a hot ratio-based analysis, we also compared the number (or rate) of false identification by making one-to-one comparison of each identification result. This pinpoint evaluation not only enables us to make a comparative analysis of each performance, but also demonstrates that our scheme more precisely identifies hot data. Lastly, we carried out experiments on variable memory size, window size, and the number of a bloom filter in our scheme to explore their impacts. Our experiments present that our proposed scheme improves the performance up to 65% although its runtime overhead of, in particular, decay operation in our scheme requires less CPU clocks up to 58% and less memory space, by an average of 50%.

References

[Agrawal 2008] Nitin Agrawal, Vijayan Prabhakaran, Ted Wobber, John Davis, Mark Manasse, and Rina Panigrahy. Design Tradeoffs for SSD Performance. In *USENIX*, 2008.

[Bauer 1995] M. Bauer, R. Alexis, G. Atwood, B. Baltar, A. Fazio, K. Frary, M. Hensel, M. Ishac, J. Javanifard, M. Landgraf, D. Leak, K. Loe, D. Mills, P. Ruby, R. Rozman, S. Sweha, S. Talreja, and K. Wojciechowski. A multilevel-cell 32 mb flash memory. pages 132–133, 351, feb. 1995.

[Bloom 1970] Burton H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Commun. ACM*, 13(7), 1970. ISSN 0001-0782.

[Boboila 2010] Simona Boboila and Peter Desnoyers. Write Endurance in Flash Drives: Measurements and Analysis. In *FAST*, 2010.

[Chang 2002] Li-Pin Chang and Tei-Wei Kuo. An Adaptive Striping Architecture for Flash Memory Storage Systems of Embedded Systems. In *RTAS*, 2002.

[Chang 2005] Li-Pin Chang and Tei-Wei Kuo. Efficient Management for Large-scale Flash Memory Storage Systems with Resource Conservation. In *ACM Transactions on Storage*, volume 1, 2005.

[Chang 2004] Li-Pin Chang, Tei-Wei Kuo, and Shi-Wu Lo. Real-time garbage collection for flash-memory storage systems of real-time embedded systems. *ACM Transactions on Embedded Computing Systems*, 3(4):837–863, 2004. ISSN 1539-9087.

[Chang 2007] Yuan Chang, Jen Hsieh, and Tei Kuo. Endurance enhancement of flash-memory storage systems: an efficient static wear leveling design. In *DAC*, 2007.

[Chiang 1997] Mei-Ling Chiang, Paul C. H. Lee, and Ruei chuan Chang. Managing flash memory in personal communication devices. In *Proceedings of the 1997 International Symposium on Consumer Electronics*, pages 177–182, 1997.

[Debnath 2009] Biplob Debnath, Sunil Subramanya, David H.C. Du, and David J. Lilja. Large Block CLOCK (LB-CLOCK): A Write Caching Algorithm for Solid State Disks. In *MASCOTS*, 2009.

[Dharmapurikar 2006] S. Dharmapurikar, P. Krishnamurthy, and D.E. Taylor. Longest prefix matching using bloom filters. *IEEE/ACM Transactions on Networking*, 14(2):397 – 409, apr. 2006. ISSN 1063-6692.

[Gupta 2009] Aayush Gupta, Youngjae Kim, and Bhuvan Urgaonkar. DFTL: a Flash Translation Layer Employing Demand-based Selective Caching of Page-level Address Mappings. In *ASPLOS*, 2009.

[Hsieh 2006] Jen-Wei Hsieh, Tei-Wei Kuo, and Li-Pin Chang. Efficient Identification of Hot Data for Flash Memory Storage Systems. *ACM Transactions on Storage*, 2(1), 2006.

[Kim 2008] Hyojun Kim and Seongjun Ahn. BPLRU: A Buffer Management Scheme for Improving Random Writes in Flash Storage. In *FAST*, 2008.

[Lee 2007] SangWon Lee, DongJoo Park, TaeSun Chung, DongHo Lee, Sangwon Park, and HaJoo Song. A Log Buffer-based Flash Translation Layer Using Fully-associative Sector Translation. *ACM Transactions on Embedded Computing Systems*, 6 (3), 2007.

[Microsoft 2007] Microsoft. SNIA IOTTA Repository: MSR Cambridge Block I/O Traces. `http://iotta.snia.org/traces/list/BlockIO`, 2007.

[Narayanan 2008] DemDushyanth Narayanan and Austin Donnelly. Write Off-Loading: Practical Power Management for Enterprise Storage. In *FAST*, 2008.

[Nath 2007] Suman Nath and Aman Kansal. FlashDB: Dynamic Self-tuning Database for NAND Flash. In *IPSN*, 2007.

[Park 2008] Chanik Park, Wonmoon Cheon, Jeonguk Kang, Kangho Roh, Wonhee Cho, and Jinsoo Kim. A reconfigurable FTL (flash translation layer) architecture for NAND flash-based applications. *ACM Transactions on Embedded Computing Systems*, 7(4), 2008.

[Shin 2009] Ji Shin, Zeng Xia, Ning Xu, Rui Gao, Xiong Cai, Seungryoul Maeng, and Feng Hsu. FTL Design Exploration in Reconfigurable High-Performance SSD for Server Applications. In *ICS*, 2009.

[Soundararajan 2010] Gokul Soundararajan, Vijayan Prabhakaran, Mahesh Balakrishnan, and Ted Wobber. Extending SSD Lifetimes with Disk-Based Write Caches. In *FAST*, 2010.

[UMass 2002] UMass. OLTP Trace from UMass Trace Repository. `http://traces.cs.umass.edu/index.php/Storage/Storage`, 2002.

[Wu 2006] Chin-Hsien Wu and Tei-Wei Kuo. An adaptive two-level management for the flash translation layer in embedded systems. In *ICCAD*, 2006.