

An In-Module Disturbance Barrier for Mitigating Write Disturbance in Phase-Change Memory

Hyokeun Lee[✉], *Member, IEEE*, Seungyong Lee, Byeongki Song, Moonsoo Kim, Seokbo Shim, Hyuk-Jae Lee[✉], *Member, IEEE*, and Hyun Kim[✉], *Member, IEEE*

Abstract—Write disturbance error (WDE) appears as a serious reliability problem preventing phase-change memory (PCM) from general commercialization, and therefore several studies have been proposed to mitigate WDEs. Verify-and-correction (VnC) eliminates WDEs by always verifying the data correctness on neighbors after programming, but incurs significant performance overhead. Encoding-based schemes mitigate WDEs by reducing the number of WDE-vulnerable data patterns; however, mitigation performance notably fluctuates with applications. Moreover, encoding-based schemes still rely on VnC-based schemes. Cache-based schemes lower WDEs by storing data in a write cache, but it requires several megabytes of SRAM to significantly mitigate WDEs. Despite the efforts of previous studies, these methods incur either significant performance or area overhead. Therefore, a new approach, which does not rely on VnC-based schemes or application data patterns, is highly necessary. Furthermore, the new approach should be transparent to processors (i.e., in-module), because the characteristic of WDEs is determined by manufacturers of PCM products. In this paper, we present an in-module disturbance barrier (IMDB) that mitigates WDEs on demand. IMDB includes a two-level hierarchy comprising two SRAM-based tables, whose entries are managed with a dedicated replacement policy that sufficiently utilizes the characteristics of WDEs. The naive implementation of the replacement policy requires hundreds of read ports on SRAM, which is infeasible in real hardware; hence, an approximate comparator is also designed. We also conduct a rigorous exploration of architecture parameters to obtain a cost-effective design. The proposed method significantly reduces WDEs without noticeable speed degradation or additional energy consumption compared to previous methods.

Index Terms—Phase-change Memory, non-volatile memory, write disturbance, in-module approach

1 INTRODUCTION

PHASE-CHANGE memory (PCM) is gaining attention as the next-generation non-volatile memory (NVM), owing to its non-volatility, low latency, and scalability [23]. In recent years, software-defined memory has been announced to utilize NVM as high-speed storage or extended memory

interchangeably [14]. In particular, in-memory databases require data to remain in memory and be accessible with low latency; hence, a high-performance database can be developed by employing PCM as a non-volatile main memory [6], [17], [19]. Moreover, products of PCM have been tested in various environments for evaluating performance and exploring their suitable applications [30], [41]. Therefore, leveraging and enhancing PCM-related technology is crucial to attaining low-latency and large-scale memory systems in the future.

Even though PCM has attractive characteristics, it is not ready to be popularized in the consumer market, because several reliability issues still exist in PCM [11], [16], [22], [24], [29], [44]. In particular, write disturbance error (WDE) is one of the major problems, which delays its widespread commercialization. WDE is an interference problem on adjacent cells similar to row-hammer in DRAM [21]. This problem must be addressed as the highest priority because it would be exacerbated as process technology shrinks [39]. Additionally, in-memory database directly store data in NVM by utilizing cache-line flushes [17], [19]. This kind of application would incur frequent write operations, making cells vulnerable to WDEs.

Previously, various approaches have been reported to mitigate WDEs in PCM devices [5], [10], [11], [16], [18], [39], [40]. Approaches based on verification-and-correction (VnC) are able to eliminate all WDEs [38], [40]. However, VnC incurs additional read operations for checking the existence of errors, degrading the performance significantly. Encoding-based schemes [10], [11], [12], [18], [36], [39] reduce the number of WDE-vulnerable data patterns with little reliance

- Hyokeun Lee, Seungyong Lee, and Hyuk-Jae Lee are with the Inter-University of Semiconductor Research Center, Department of Electrical and Computer Engineering, Seoul National University, Seoul 08826, South Korea. E-mail: {hkleee, sylee, hyuk_jae_lee}@capp.snu.ac.kr.
- Byeongki Song and Moonsoo Kim are with Samsung Inc., Hwasung, Gyeonggi-do 18448, South Korea. E-mail: bksong@capp.snu.ac.kr, ms213.kim@samsung.com.
- Seokbo Shim is with SK Hynix Inc., Icheon, Gyeonggi-do 17336, South Korea. E-mail: seokbo.shim@sk.com.
- Hyun Kim is with the Department of Electrical and Information Engineering and the Research Center for Electrical and Information Technology, Seoul National University of Science and Technology, Seoul 01811, South Korea. E-mail: hyunkim@seoultech.ac.kr.

Manuscript received 4 November 2021; revised 21 June 2022; accepted 23 July 2022. Date of publication 8 August 2022; date of current version 13 March 2023.

This work was supported in part by the National Research Foundation of Korea(NRF) grant funded in part by the Korea government(MSIT) 2022R1F1A1062786, in part by the (MSIT) Ministry of Science and ICT, Korea, under the ITRC (Information Technology Research Center) support program (IITP-2022-2020-0-01461, in part by the IITP Institute for Information & communications Technology Planning & Evaluation, in part by the Technology Innovation Program 20011074, Development of Open Convergence Memory Solution and Platform for Next Generation Memories, and in part by the Ministry of Trade, Industry & Energy(MOTIE, Korea. The EDA tool was supported by IC Design Education Center (IDEC), Korea.

(Corresponding author: Hyun Kim.)

Recommended for acceptance by M. T. Kandemir.

Digital Object Identifier no. 10.1109/TC.2022.3197071

on VnC, but the mitigation performance of these approaches varies considerably with data patterns in applications. Studies [16], [34] have reported that WDEs may occur when a cell experiences more than a specific number of RESET pulses from its neighbors, which is more realistic than a random WDE model. Although the study in [16] has presented the manufacturing metric that incurs WDEs, their approach leverages a write-cache to reduce the write traffic without considering such a realistic WDE model. Furthermore, a large capacity of SRAM is required for mitigating WDEs notably. For above reasons, given that previous approaches are entirely decoupled from this realistic model, new approaches that manage aggressors, which are actively programmed cells that likely incur WDEs on neighboring cells, are necessary with negligible performance overhead in PCM modules.

To satisfy these requirements, this paper proposes an in-module write disturbance barrier (IMDB) that utilizes a realistic WDE model and restores vulnerable data on demand. Because the realistic WDE model shows that WDEs occur with a specific number of neighboring writes, the proposed method records the number of RESETs in a table. Using the recorded information, most of the WDE-vulnerable data can be rewritten before the occurrences of WDEs, and only addresses need to be managed in the data structure to reduce the burden on the supercapacitors upon system failure. For further error mitigation, a tiny data cache, referred to as a *barrier buffer*, is introduced to store highly aggressive address information. Meanwhile, the replacement policy may expand the number of read ports on SRAM, involving a considerable overhead. This is because the policy merely regards the entry holding a smaller number of 1-to-0 flips as an eviction candidate. Therefore, an approximate lowest number estimator (*AppLE*), which probabilistically counts the numbers based on the sampling method, is proposed to accommodate the use of a dual-port SRAM (DPSRAM) without speed degradation. Experimental results indicate that our approach reduces WDEs compared to previous studies, with negligible overhead. In conclusion, the contributions of this study can be summarized as follows:

- The first on-demand WDE mitigation method is proposed. Based on a more practical WDE trigger model, the proposed method leverages a two-level SRAM and restores vulnerable cells on demand.
- This paper introduces a novel prior-knowledge-offering method, because the replacement policy may contradict the locality of applications.
- The replacement policy requires hundreds of ports on an SRAM in a naive approach. This paper designs probabilistic hardware, *AppLE*, to allow the use of a DPSRAM for enhancing the feasibility.
- Several design parameters are required in the proposed method; hence, rigorous sensitivity analyses are conducted to acquire the cost-effective design.

2 BACKGROUND AND MOTIVATION

2.1 Introduction to Phase-Change Memory

PCM is a non-volatile memory device that has two different states, *amorphous* and *crystalline*. The former has a higher resistance than the latter [24]. The detailed overview of a

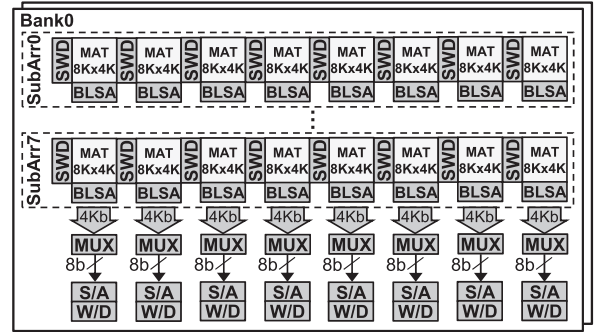


Fig. 1. Architecture of a PCM device.

PCM device in an 8GB dual-rank module is illustrated in Fig. 1. The device consists of eight subarrays, and each subarray is composed of eight cell matrices (MATs). Main wordline drivers activate a subarray in each bank. Using the row address, each sub-wordline driver (SWD) activates 4Kb data. The activated data are sensed by bitline sense amplifiers (BSA) and transferred through global bitlines. Using the column address, each column multiplexer (MUX) outputs an 8-bit word to global sense amplifiers (S/A) by multiplexing 4Kb data. Finally, 8 words are transferred to the data bus in burst mode. In total, 64B are carried out from eight devices, which are driven symmetrically by a single command. For a write operation, data on write drivers (W/D) are written back to the cell array.

2.2 Modeling Write Disturbance in PCM

WDE is caused by the resistance shift from the amorphous state to the crystalline state [18], [34], [40]. WDEs occur on an idle cell adjacent to the cell under RESET operations [18], [40]. Since the intensity of current during a SET operation is nearly half of that during a RESET operation, an idle cell's temperature next to the programmed cell would be higher than those under SET (but lower than those cells under RESET). As a consequence, a phase transition may occur on that idle cell.

Knowing the occurrence moment of WDEs is also crucial for modeling WDEs in a simulator. Rather than triggering WDEs randomly, the study in [34] explains when a WDE occurs, according to low-level characteristics of WDE. It shows that an amorphous cell gradually shifts to crystalline state due to heat transfer to neighbors, thereby incurring WDEs. The study also explains that a cell can be programmed in different time frames using one pulse per frame; hence, WDE can occur regardless of the idle time duration between consecutive writes. In this paper, we refer the number of pulses incurring WDEs as the *WDE limitation number*.

The prior work in [16] reports the WDE limitation number as 5K-10K, but the author in [16] does not set that number for evaluation. Instead, both this study and [16] assume that a WDE occurs when the number of writes (i.e., 1-to-0 bit flips) exceeds the WDE limitation number of 1K, and this is uniformly applied to all cells. This is because setting the number of 5K or 10K requires a much longer simulation time for triggering WDEs in a row. Our proposed method can be simply extended to various WDE limitation numbers, because the threshold for generating rewrite command is formalized as the function of the WDE limitation number.

TABLE 1
Performance of Randomized VnC

Probabilities of VnC			WDE reduction	Speedup
both rows	upper row	lower row		
0%	50%	50%	23%	57%
75%	12.5%	12.5%	30%	18%
80%	10%	10%	36%	17%
90%	5%	5%	36%	15%
95%	2.5%	2.5%	43%	15%
99%	0.5%	0.5%	46%	14%

Furthermore, the industry has presented that WDEs mainly occur on adjacent materials patterned on a common bitline [25]. This is because PCM cells are overlapped with bitlines, incurring simpler heat dissipation along bitlines. Therefore, WDEs mainly occur on adjacent materials patterned on a common bitline. However, our proposed scheme can be easily extended when more than two neighbor cells are disturbed by generating more rewrite operations, which are used for restoring vulnerable cells on demand.

2.3 Motivation

Necessity of Reducing the Cache Burden. Cache-based schemes mitigate WDEs by temporarily storing write data into dedicated SRAM. Although a cache-based scheme (i.e., SIWC [16]) can significantly reduce the number of WDEs in PCM compared to those in previous studies (see Section 5.7), this strategy requires high-capacity SRAM, because it indiscriminately caches write data. Furthermore, data adjacent to cached addresses remain vulnerable to WDEs. To overcome these challenges, it is necessary to store the data that likely incur WDEs (i.e., WDE aggressors) and restore cells adjacent to these aggressors. Therefore, a preprocessor that filters non-aggressors and restores cells adjacent to aggressors is necessary (called “main table”) if a small-sized cache (called “barrier buffer”) is desirable.

Necessity of Reducing the Performance Overhead of VnC. VnC, the most common solution to WDEs, triggers read operations to read two neighboring data before the objective data is updated. Subsequently, two neighbors are read again after the write operation for verification. Finally, VnC is performed iteratively if WDEs occur on the neighbors, degrading the performance markedly by these read operations. A naive approach to reducing the number of such read commands is to perform VnC randomly. Table 1 shows the WDE reduction rate in comparison with the baseline (i.e., no VnC) and the speedup in comparison with the normal VnC (i.e., always verify both rows). For example, the third row assumes probabilities of this tuple are 80%, 10%, and 10%, respectively. Random VnC yields a 14% speedup compared to normal VnC and a WDE reduction rate of 46% compared to the baseline. This is because PCM does not require a refresh operation by default (or an infrequent refresh compared to DRAM), causing cells scarcely to be restored. In contrast, high speedup (i.e., 57%) is attainable at the expense of reliability. Moreover, the operations of VnC (i.e., pre-write read, write, and post-write read) are strictly ordered; hence, the speedup is not notable even when a probabilistic approach is applied. Please note that these data are extracted based on the configuration in Section 5.1. As a result, the

TABLE 2
Characteristics of Representative Schemes

Schemes	LAZY [40]	ADAM [39]	SIWC [16]	IMDB
Approach	VnC	Encode	Cache	Demand
WDE reduction	High	Low	Moderate	High
Speed	Low	Moderate	Moderate	Moderate
Energy	High	Moderate	Moderate	Moderate
Storage	Very large	-	Large	Small

VnC-based scheme is unsuitable as a preprocessor (i.e., main table) for the filtering mentioned above. Thus, there is need for a new on-demand approach that accurately predicts vulnerable patterns and reduces the number of WDEs to a small value comparable to VnC.

Table 2 shows the relative characteristics of previous WDE mitigation schemes (explained in Section 7) against our proposed method, IMDB. The VnC-based approach (i.e., LAZY or lazy correction) incurs a significant performance and energy overhead due to the increased number of commands for verification. LAZY requires an additional “WDE-free” error correction pointer (ECP) device with a lower density than the normal device [40]. ADAM only requires simple compression logic without storage resources; however, the mitigation performance is much lower than IMDB due to high dependency on application data patterns. SIWC reduces WDEs moderately by introducing a write cache, which is larger than that of IMDB. Meanwhile, IMDB significantly reduces WDEs to the number similar to the lazy correction by restoring vulnerable data on demand with a small SRAM.

3 IMDB: IN-MODULE DISTURBANCE BARRIER

3.1 Architectural Overview

Fig. 2 depicts the overall architecture, where NVM commands are dispatched from the integrated memory controller in the host. For the PCM module, the media controller generates micro-commands and schedules commands to available banks in the media devices. A DRAM cache is only used for storing an address indirection table (AIT) [15], [43]. The proposed module, IMDB, is located between the media controller and media devices.

As shown in Fig. 2, IMDB consists of the main table (Section 3.2.1), a barrier buffer (Section 3.2.2), and AppLE (Section 4.2). First, the main table manages the addresses of WDE aggressors. If a write address hits in the table, the number of 1-to-0 bit flips is calculated and accumulated in the table; otherwise, the dedicated replacement policy supported by AppLE, which reduces the overhead incurred by multi-port SRAM, selects a victim entry within the table and replaces it with the new address. When the number of bit flips on the aggressor exceeds the pre-defined threshold, IMDB generates *rewrite commands* for data that are adjacent to the aggressor. As explained in Section 2.2, an idle cell in amorphous state (i.e., RESET) gradually shifts to crystalline state if it is exposed to high-temperature several times. Then, a WDE happens when this cell completely turns into crystalline state. Therefore, the rewrite command is introduced and




Fig. 2. Architectural overview of the proposed system. Please note that each PCM media device follows the architecture layout of Fig. 1.

used for restoring such partially shifted cells back to amorphous states before the occurrences of WDEs. Subsequently, IMDB migrates the information from the main table to the barrier buffer that comprises a few data entries, reducing WDEs further. Even though the bit width of a barrier buffer's entry is longer than that of the main table, the barrier buffer manages much fewer entries; hence, it occupies less SRAM capacity than the main table. Fig. 2 shows the swapping mechanism between two tables, by which WDE aggressors are managed as long as possible within IMDB.

Our proposed work, IMDB, is a new approach to mitigating WDEs. In particular, IMDB differs from wear leveling and previous WDE studies. Wear leveling uniformizes the number of write accesses across different physical regions; however, it just temporally defers WDEs. In contrast, IMDB estimates WDE-vulnerable addresses by utilizing the WDE limitation number and recording aggressors. The estimated vulnerable addresses are then restored to stable states. Furthermore, it is noteworthy that the wear leveling does not affect the threshold selection because the threshold for generating rewrite commands is derived from WDE limitation number, which is determined by the cell characteristics. Indeed, the wear leveling spreads the number of writes over all PCM regions, making it possible to lower occurrences of WDEs *within a fixed time interval*. However, the wear leveling cannot increase the threshold for generating rewrite commands. This is because the wear leveling just temporally postpones occurrences of WDEs. For example, when a data is remapped from cell-A to cell-B due to wear leveling, the state of cell-A remains shifted (i.e., between amorphous and crystalline) because PCM does not require erase operations. Thus, cell-A is still vulnerable to WDEs if another data is mapped to cell-A. It should be noted that WDEs depend on the number of 1-to-0 bit flips on neighboring cells regardless of the rate of programming pulses, as explained in Section 2.2. Consequently, wear leveling is an orthogonal methodology compared to IMDB; wear leveling only defers WDEs rather than reducing WDEs. Furthermore, one of recent studies related to WDEs shows that simply remapping data (e.g., start-gap[33] or security-refresh [35]) has small effects on reducing WDEs [20]. On the other hand, IMDB reduces occurrences of WDEs by directly estimating WDE-vulnerable addresses in the main table and barrier buffer.

In general, there are three categories for mitigating WDEs: VnC-based schemes [38], [40], encoding schemes [10], [12], [36], [39], and the cache-based scheme [16]. The VnC-based method defers correction by assuming no error in the additional device; however, VnC basically incurs high performance overhead. In contrast, IMDB restores data

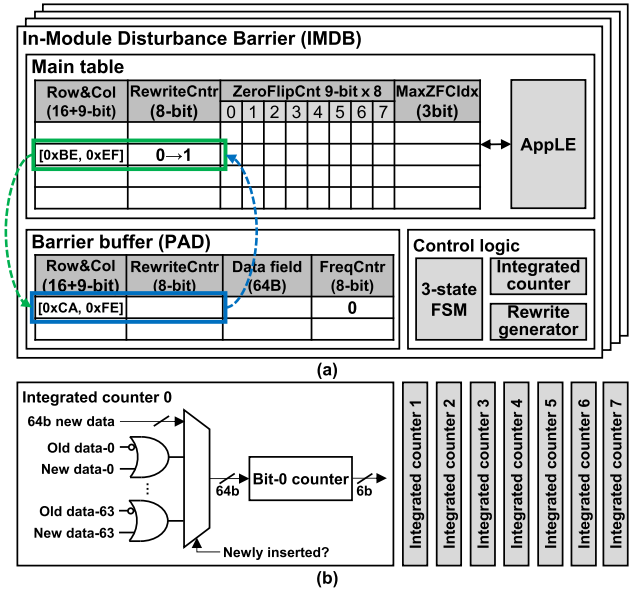


Fig. 3. Detailed design of IMDB: (a) implementation of four IMDB planes (each IMDB plane is assigned to each PCM bank operation), (b) integrated counters for eight *ZeroFlipCnts*.

before WDEs occur. Encoding schemes are highly dependent on application data patterns. Compared with this kind of schemes, IMDB monitors the vulnerability of data patterns, leading to less dependence on application data pattern. The cache-based scheme requires a larger SRAM for notably reducing WDEs. On the other hand, IMDB buffers urgent data using the WDE limitation number, reducing WDEs significantly with an SRAM capacity that is four times smaller than the previous study.

3.2 Implementation of Data Structures

Fig. 3a shows the detailed architecture of IMDB, where each plane is allocated for every PCM bank; hence, all IMDB planes operate concurrently at the bank level without contention. An IMDB plane consists of two tables, namely a *main table* and a *barrier buffer*. The following subsections describe implementations of each table.

3.2.1 Main Table

The main table is implemented with a set of SRAMs, where the entry is updated by a control logic. In particular, four fields in the table are used for estimating the degree of WDE on write addresses:

- *Row & Col*: Indicates row and column addresses that are currently being managed.
- *ZeroFlipCntr*: Eight sub-counters are in the field, each of which counts the number of bit flips from 1 to 0 and manages one 64-bit word in a 64B cache line. Each of the eight *ZeroFlipCnts* manages a 64-bit word within a device (or chip), because one 64-bit word outputs from each of eight devices, as shown in Fig. 1. Consequently, these eight *ZeroFlipCnts* map to a row of a subarray and symmetrically manage eight 8-bit sub-words across eight devices.
- *MaxZFCIdx*: Indicates the sub-counter index of *ZeroFlipCntr* holding the maximum value. It is updated

in control logic after reading an entry. It is used for comparing the maximum value of the *ZeroFlipCntr* with the threshold value for rewrite operations.

- *RewriteCntr*: Represents the frequency of rewrite operations on the address of *Row & Col* in an 8-bit counter.

A per-bank IMDB plane is assigned to each bank; hence, bank parallelism is ensured to lower the contention on IMDB. Furthermore, IMDB prevents resource redundancy, because only one command processed in IMDB at a time without incorporating a serialized queue. The command is handled by a three-state finite state machine (i.e., IDLE, HIT, MISS) in control logic, where the varying latency of the multiple states are factored in the simulator. After a command is inserted, IMDB operates in two different ways, depending on the existence of the address in the table:

- If the address is found in the main table, the state transits to HIT. Meanwhile, two types of data, i.e., the new write-data and the previously written data already read in the controller, are passed to control logic. Subsequently, the number of 1-to-0 bit flips is counted by integrated counters (see Section 4.2) and accumulated to the corresponding *ZeroFlipCntr*. When the maximum value of *ZeroFlipCntr* surpasses the predefined *threshold*, two rewrites on adjacent wordlines are generated and sent to the write queue in the media controller. Accordingly, the value of *RewriteCntr* increases.
- If an address is not found in the main table, an insertion is required while converting the state to MISS. The probabilistic insertion method is leveraged in this study, where infrequent accesses are filtered out with probability p to reduce evictions from the SRAM. When insertion is required, our proposed replacement policy determines the victim (explained in Section 4), and thereby the new address can replace the victim entry.

According to hit/miss cases on the main table, the finite state machine is a trigger for different operations. For both cases, after the table reference, write data issued to the media right away. Memory commands in the media controller scheduler must follow the promised timing constraints. Thus, no command can be entered to the same IMDB plane during the write phase in the media, allowing the background processing of IMDB.

In the proposed design, two parameters, (1) the threshold of generating rewrite commands and (2) the probability p , are necessary. First of all, we decide the threshold of generating rewrite commands in the main table as *WDE limitation number/2-1*, because two rows can disturb a row. Thus, if we assume a WDE limitation number of 1K, as in [16], the threshold becomes 511, making the bit width of each *ZeroFlipCntr* to be 9. The other parameter, p , indicates the probability of inserting a new missed address into the main table. Increasing the probability incurs more frequent entry replacement in the table for detecting WDE aggressor, losing the opportunity to rewrite the victims of WDEs. In contrast, lowering the probability makes “long-term” attacks lose the chance to be in the table. Our experiments regarding different insertion probabilities show that $p=1/128$ yields the fewest WDEs; hence, we select $p=1/128$.

As shown in Fig. 3a, the main table employs two types of SRAMs. First, a dual-port content-addressable SRAM (CAM) is allocated as *Row & Col* fields. Second, a multi-port SRAM, consisting of *ZeroFlipCntr*, *MaxZFCIdx*, and *RewriteCntr*, has multiple read ports for obtaining all entry contents at once to apply the proposed replacement policy (see Section 4.1). However, since the use of multi-port SRAMs causes a significant overhead, we propose *AppLE*, which enables the replacement policy with a DPSRAM without speed degradation (see Section 4.2).

3.2.2 Barrier Buffer

The barrier buffer is introduced to store the data with frequent 1-to-0 bit flips. For a read request, the barrier buffer is capable of serving commands directly. For a write command, if the address hits on the barrier buffer, the data are updated in the barrier buffer directly. Otherwise, if an address hits only on the main table, the normal operation of the main table is performed, as explained in Section 3.2.1.

As shown in Fig. 3a, the green-boxed entry in the main table contains the data frequently exposed to 1-to-0 flips. It is invalidated and *promoted* to the barrier buffer when *RewriteCntr* updates (i.e., rewrite occurs in the main table). The barrier buffer inherits the address and *RewriteCntr* information from the main table. If the barrier buffer is not full, the promoted entry can be directly placed in the barrier buffer. After several entry promotions (i.e., rewrite operations) from the main table, the barrier buffer would become full. At this moment, the promoted entry (from the main table) replaces the least frequently used (LFU) entry that is bounded by the blue box in Fig. 3a. For this reason, *FreqCntr* is required for the replacement policy, as in [32]. The LFU entry data are then sent back to the media controller for writing back the dirty data, and this information is demoted to the main table. Because the demoted addresses have been WDE aggressors before, the number of rewrites is reserved in *RewriteCntr*. *RewriteCntr* provides historical information with which to obtain a reasonable victim candidate in the main table (explained in Section 4.1). Please note that the 8-bit of *RewriteCntr* is a generously selected bit width to prevent overflow based on our experiments.

To implement the barrier buffer, a dual-port CAM-based SRAM and a dual-port SRAM are employed for *Row & Col* and *data & RewriteCntr & FreqCntr*, respectively. The energy consumption is negligible, because only a small number of entries in the barrier buffer are necessary to provide high WDE mitigation performance, as shown in Section 5.7.3. The sensitivity analysis of the number of entries will be shown in Section 5.5.

3.3 Modification of Media Controller

The media controller is modified to support IMDB in two aspects. First, acquiring the old data is necessary to count bit flips. Thus, a *pre-write read operation* is performed ahead of a write command. The pre-write read request has a higher priority than write requests but a lower priority than normal read requests because write requests in the controller mainly drain when the queue is full. Lastly, a merge operation is introduced, by which the rewrite command can coalesce with a same-address write command.




Fig. 4. A toy example showing malicious attacks. 0xDEAD evicts insufficiently baked 0xBEEF, which is vulnerable to WDEs with gradual 1-to-0 bit flips.

4 REPLACEMENT POLICY

4.1 Replacement Policy for IMDB

A replacement (or eviction) policy is required in the main table based on the characteristics of WDEs. Therefore, we exploit *ZeroFlipCntr* and *RewriteCntr* to define the replacement policy. When the input command requests a new entry in the main table, the policy is ready to select the victim entry. The victim candidate is defined as a less urgent aggressor, thereby selecting the minimum value of *ZeroFlipCntr*. However, more than two candidates may exist if the table has multiple entries with the same values of *ZeroFlipCntr*. Since the aggressiveness of WDEs varies with historical information (i.e., *RewriteCntr*), the entry containing the minimum of *RewriteCntr* is finally selected as the replaced entry.

To prevent "cold-start" that incurs early eviction from the table, this study introduces *prior knowledge*. Since the policy prioritizes the present vulnerability using *ZeroFlipCntr*, the recently inserted but insufficiently "baked" entry can easily be evicted from the main table. Although *RewriteCntr* contains the historical information, it becomes useless if the entry is newly inserted and evicted right away (see example in Fig. 4). To tackle this problem, the prior knowledge, which is simply defined as the number of zeros in each data block, is stored in *ZeroFlipCntr*.

It is noteworthy that a module, namely *integrated counter*, is required to perform the above processes. The integrated counter provides mainly two functions. First, it counts the number of 0s of newly inserted data, which is then directly used as *prior knowledge* of *ZeroFlipCntr*. Second, it counts the number of 1-to-0 bit flips of the accessed address in the table. The counted value is then added to the *ZeroFlipCntr*. As a result, the integrated counter is implemented as Fig. 3b, where eight counter blocks are required to count each 64-bit word in a 64-byte data concurrently.

4.2 Approximate Lowest Number Estimator

The eviction policy requires the number of read ports to be equal to the number of entries on the main table. It increases latency, area, and energy overheads. If a 256-entry main table is assumed, 255 tree-structured dual-input comparators are necessary for latency minimization (i.e., 8 cycles). However, our evaluation results in Figs. 5a and 5b indicate that increasing the number of read ports on an SRAM significantly increase overheads. As a result, an SRAM with 256 read ports is an infeasible implementation.

To reduce such overheads, this paper introduces a sampling-based comparator, called AppLE. The basic concept of AppLE is to bind multiple entries. For example, binding 8 entries results in 32 groups. In this case, a randomly generated number ranging from 0 to 7 is multiplied by 8 and assigned to




Fig. 5. Characteristics of a 256-entry SRAM having multiple read ports, which is extracted from CACTI [9]: (a) energy, (b) latency and area.

each group (i.e., $group_index \times 8$). This assigned value is used as the main table's input address, and a *sampled entry* is referenced. Using this addressing mechanism, the victim candidate is selected among sampled entries.

The main concept of AppLE is comparing counts approximately by grouping a few entries in the main table, instead of comparing counts in parallel. Two design options are first discussed in this paper: the first one naively implements approximate counting in parallel (Fig. 6a); the second one performs approximate counting sequentially without increasing the latency on the critical path (Fig. 6b). The first option (i.e., the parallel one) is infeasible to be implemented in the industry, because it simply regards the number of groups as the number of read ports. For example, the typical I/O frequency of DDR4 is around 800MHz [26], and the maximum target number of read ports is set to 32. Still, the area of a 32-port SRAM is $105\times$ larger than that of a single read port SRAM. Moreover, an SRAM consisting of dozens of read ports is unusual in terms of manufacturing.

This is the reason for choosing the second design option. In the second design, the latency of sequential comparisons can be hidden within the IDLE state. This is because the number of comparisons reduces with AppLE (e.g., 32 cycles for above example), and the IDLE state maintains for 120 cycles after issuing a write command. We directly evaluate the case of comparing all 256 entries (i.e., no-AppLE) in Fig. 12b; it shows no-AppLE case incurs 15% of performance degradation, because 136 cycles (=256-120) of additional latency cannot be hidden within the IDLE state.

5 EVALUATIONS

5.1 Configurations

Table 3 shows the configuration of evaluation environment. In this study, we use four simulators to simulate a PCM-based

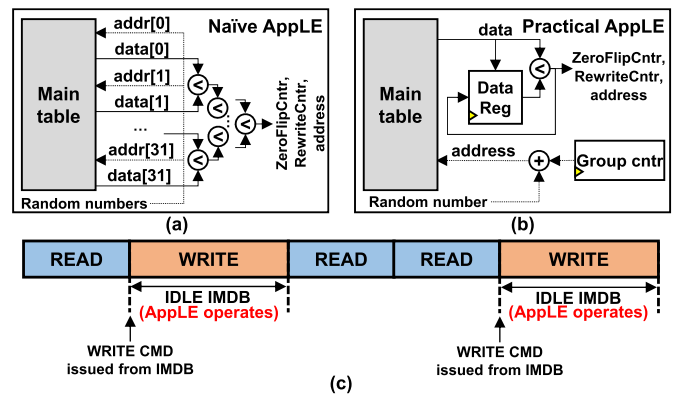


Fig. 6. Implementations of AppLE: (a) a naive approach, (b) a practical approach, (c) timeline of IDLE state in IMDB.

TABLE 3
Simulation Configurations

Simulator	Device	Description
gem5	Cores	Out-of-order, 4-core, 2GHz
	L1 cache	I-cache: 2-way set associative, D-cache: 4-way set associative, each has a capacity of 64KB.
	L2 cache	Shared last-level cache. 16-way set associative, 1MB.
NVMain	Media controller	Separated write queue and read queue (64-entry), FR-FCFS.
	PCM	Read: 100ns, RESET: 100ns, SET: 150ns Write disturbance limitation: 1K Size: 8GB (2-rank, 2-bank/rank)

main memory system: gem5[2], NVMain[31], NVSim[7], and CACTI[9]. Gem5 is a processor architecture simulator that is configured as a quad-core processor [2]. NVMain is a simulator that simulates details of NVM subsystems [31]. Both simulators are functional- and cycle-accurate; hence, running gem5 and NVMain together requires extremely long simulation time. Moreover, sensitivity analysis requires more than 400 experiments in this paper. Thus, trace-driven simulation is necessary to significantly reduce the simulation time. Trace-driven simulation is a common evaluation methodology in NVM-related studies, as performed in [10] and [39]. To conduct the trace-driven simulation, we first extract memory command traces by running workloads on gem5 in standalone mode. Thereafter, extracted command traces are fed into NVMain, which can also be run in a standalone manner. NVSim [7] and CACTI [9] are energy simulators to estimate energy parameters (i.e., energy per access) of PCM and SRAM. The energy evaluation mechanism in NVMain calculates the energy consumption of two memory types using energy parameters obtained from these two energy simulators. Still, a large L2 cache in the processor requires a long simulation time to incur enough WDEs (i.e., more than 100); hence, it is necessary to determine a small but practical L2 cache size to build a burn-in test environment. Therefore, the processor is configured as the mobile processor [1], which may incur increased memory traffic. Nonetheless, it should be noted that we extract memory traces having a wide range of misses per thousand instructions (MPKI) in order to simulate the various kinds of memory traffic, as shown in Table 4. In this study, traces are obtained from SPEC CPU benchmark suit [8] and synthesized persistent workloads (prefixed as “pmix”) that are similar to those in [6], [17], [19]. Please note that the *baseline* does not apply any mitigation scheme.

5.2 Architectural Exploration

Design parameters, specifically the number of entries in the main table (N_{mt}), the number of entries in the barrier buffer (N_b), and the group size dedicated to AppLE (N_g), are crucial when seeking a cost-effective architecture for IMDB. As explained in the previous section, the latency of AppLE can be entirely hidden by the IDLE state of IMDB from $N_g = 32$ (see Fig. 5), which also holds for $N_g < 32$. Moreover, 64 is determined as the maximum number of entries in the barrier buffer to guarantee that no more than 10% of the flush time (i.e., 100us) is consumed. As a result, the trade-off function of IMDB is defined as follows:

TABLE 4
Information on Workloads

Workloads	Description	MPKI
SPEC::bzip2	General compression	11.98
SPEC::sjeng	Artificial intelligence (chess)	0.89
SPEC::h264ref	Video compression	1.65
SPEC::gromacs	Biochemistry	5.49
SPEC::gobmk	Artificial intelligence (go)	6.65
SPEC::namd	Biology	1.09
SPEC::omnetpp	Discrete event simulation program	6.99
SPEC::soplex	Linear programming optimization	21.31
pmix1	Queue, Hashmap, B-tree, Skiplist	10.24
pmix2	Queue, B-tree, RB-tree, Skiplist	11.10
pmix3	Hashmap, RB-tree, Queue, Skiplist	8.95
pmix4	RB-tree, Hashmap, B-tree, Skiplist	10.12

$$T = W(N_{mt}, N_b, N_g) + A(N_{mt}, N_b) + S^{-1}(N_b),$$

where $N_g \leq 32, N_b \leq 64$ (1)

where W , A , and S are the number of WDEs, the area, and the speedup (i.e., execution time normalized to the baseline [40]), respectively. Based on Eq. (1), this section evaluates the effectiveness of the prior knowledge and determines the main table size (N_{mt}). Subsequently, sensitivity analyses concerning the number of entries in the barrier buffer (N_b) and the group size for AppLE (N_g) are conducted to determine the cost-effective parameters. Finally, these parameters are applied and compared to previous studies.

5.3 Effectiveness of the Replacement Policy

Several replacement policies have been published in previous studies, such as MRU (most-recently used), LFU (least frequently used), and LRU-alike policies (e.g., pseudo-LRU). It is noteworthy that WDEs occur when a neighboring cell is frequently programmed. Thus, we need to consider this characteristic when choosing the appropriate policy. MRU discards the most recently used items. However, WDEs may occur on some applications with relatively high locality. For LFU, we need to add additional metadata on the entry for representing the access frequency, incurring higher resource costs. As a result, we finally compare the proposed policy against the LRU, because LRU simultaneously considers the locality and the access frequency. Fig. 7a shows that the LRU yields higher WDEs than the proposed policy, because the LRU makes the address close to WDEs be evicted if it is not accessed for a long time. For example, *bzip2*, *gobmk*, *gromacs*, and persistent workloads have this kind of access pattern, increasing WDEs. In contrast, the proposed policy observes the number of bit flips and keeps track of their long-term history. However, the LRU shows $3 \times$ fewer WDEs than the proposed policy on *namd*. This is because *namd* has high spatial and temporal locality. We find that *namd* achieves a 70% higher row buffer hit rate than an application of a similar MPKI (i.e., *sjeng*), yielding lower hit rate on the main table. However, such a degradation will be mitigated in the following subsections.

Fig. 7b shows that replacement policies for the main table do not affect the speed performance, because the main functionality of the IMDB is managing WDE aggressors without caching plenty of data in SRAM. On the other hand, the

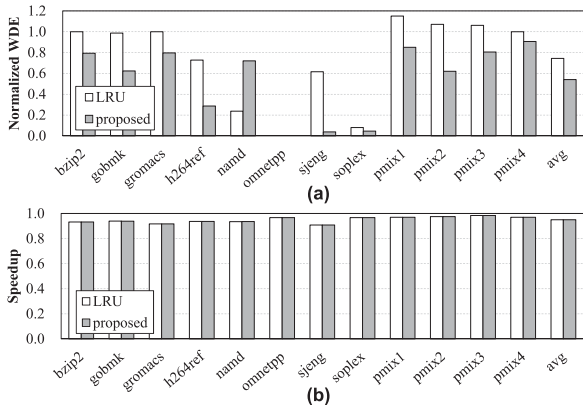


Fig. 7. Performance according to different replacement policies: (a) normalized WDE, (b) speedup.

proposed policy generally contributes to lower WDEs (see Fig. 7a), because it keeps aggressors more precisely than the LRU and rewrites rows adjacent to aggressors.

5.4 Sensitivity to Main Table Configuration

Figs. 8a and 8b show the normalized WDE regarding different numbers of entries in the main table. Both figures show that WDEs generally decrease as the number of entries increases. In particular, as shown in Fig. 8a, while the number of WDEs exceeds that in the baseline when the number of entries is fewer than 256, the number decreases sharply from 2048 entries. This is because the small-size table cannot be trained due to frequent entry replacement on the main table. On the other hand, as shown in Fig. 8b, the 256-entry main table with prior knowledge yields a result equivalent to that of the 2048-entry table without prior knowledge. In other




Fig. 9. Relationship between WPKI and the number of main table entries.

words, the proposed method yields an eightfold increase in the efficiency of the WDE mitigation performance.

Fig. 8c presents the average normalized WDE and the capacity required for the main table, and the probabilistic insertion scheme discussed in Section 3.2.1 is already adopted for both configurations. As shown in Fig. 8c, the normalized WDE is 95% lower than the case without prior knowledge at 256 entries. Furthermore, the main table's capacity significantly increases from 512 entries; hence, 256 entries can be selected as an appropriate number of entries in the main table, considering the trade-off between the performance and the area. In summary, from this subsection, the number of entries in the main table is fixed as $N_{mt} = 256$.

Rather than write request rates or write patterns (e.g., stride or stream), the number of WDEs fundamentally relies on the number of 1-to-0 bit flips on neighboring addresses. In other words, data programming patterns of applications determine the overall WDE occurrences in a PCM device. However, we find that the number of writes per thousand instructions (WPKI) determines the number of main table entries. Fig. 9 shows the relationship between WPKI and the number of main table entries; the number of table entries in this figure denotes the number of entries to fully eliminate WDEs only using the main table. In general, fewer entries are required for smaller WPKI values because WPKI determines the footprint of write commands. Moreover, a larger footprint of write commands incurs more frequent replacements on the main table, yielding more WDEs. Therefore, the main table needs to operate with the barrier buffer and AppLE for higher mitigation performance, because each component of IMDB is complementary to each other.

5.5 Sensitivity to Barrier Buffer Size

Fig. 10a shows the number of WDEs with different numbers of entries (i.e., different sizes) in the barrier buffer. For clarity, the results are normalized to the *temporal base condition*; that is, the main table consists of 256 entries with the prior knowledge. Please note that Fig. 10a only shows benchmarks still having WDEs under the temporal base condition. As shown in this figure, most benchmarks yield significantly fewer WDEs with the 4-entry barrier buffer. On the other hand, WDEs in *gobmk* decrease when the 64-entry is applied, because some write patterns have extremely long period; these are unlikely to be affected by the proposed policy regardless of the buffer size. However, the following subsection shows that AppLE resolves this problem.

Fig. 10b shows the average normalized WDE of the benchmarks mentioned above. Because the speedup does not increase remarkably considering the number of entries, S^{-1} in Eq. (1) can be referred to as a constant. Furthermore, the capacity of the barrier buffer is at least three times as small as the main table for $N_b \leq 16$ (see bit widths of tables

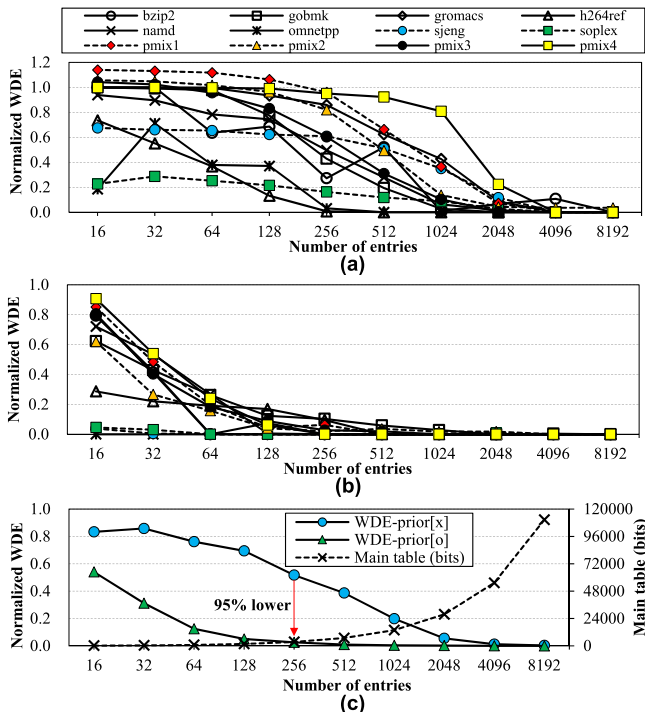


Fig. 8. Normalized WDEs regarding different numbers of entries in the main table, (a) without prior knowledge, (b) with prior knowledge, (c) displaying average normalized WDE and SRAM capacity.




Fig. 10. Sensitivity to the number of entries in barrier buffer: (a) normalized WDE, (b) average performance.

in Fig. 3a), which makes the capacity of the barrier buffer negligible compared to the main table. It can be concluded that W in Eq. (1) is sufficient to obtain a cost-effective architecture. Therefore, we select $N_b = 8$ as the trade-off point, because the WDE stabilizes from 8 entries (i.e., 76.5%).

5.6 Sensitivity to AppLE Group Size

Fig. 11a presents the absolute number of WDEs with different numbers of groups. Here, the barrier buffer is not applied for straightforward analysis, and 256 groups mean that AppLE is not applied. As presented in Fig. 11a, WDEs lower with fewer groups for most benchmarks. Furthermore, AppLE has the potential for avoiding “tricky patterns”. The worst-case behavior for WDEs can be caused by repetitive 0 and 1 pulses on the same address, which incurs WDEs on $512 \times 2 = 1024$ bits. However, the main table can easily detect such a pattern, because it manages the number of 1-to-0 bit flips and generates rewrite operations on vulnerable addresses. In contrast, a trickier way to induce WDEs is incurring 1-to-0 bit flips on an address (say “A”) with a long period (e.g., *gobmk*). Furthermore, a large number of unrepeated addresses except “A” are programmed in this long period (i.e., ABC...ADE...A...). This tricky pattern confuses the main table and frequently replaces entries; however, AppLE binds multiple entries as a group, and only one entry randomly becomes a replacement candidate within a group. Therefore, the adversarial address rarely gets evicted from the table for a larger group size. The graph of *gobmk* in Fig. 11a shows that the group size of 8 (whereby the number of groups is 32) yields lower WDEs than the case without AppLE. However, WDEs increase significantly from 2 groups (see red graph in Fig. 11a). In particular, the *fully randomized replacement policy* (i.e., one group) shows $15 \times$ more WDEs than the case without AppLE, indicating that the fully randomized replacement policy is less reliable. As a result, $N_g = 8$ or 4 is selected as an appropriate design parameter for AppLE.

Fig. 11b presents speedups regarding different numbers of groups. If AppLE is not applied, a 256-cycle latency is induced at least. Even though the latency can be hidden within the write latency (i.e., 120 cycles), at least 136 remaining cycles slow down the performance by 15%, as shown in the figure. In contrast, AppLE has no performance degradation due to latency hiding. Regarding energy consumption,

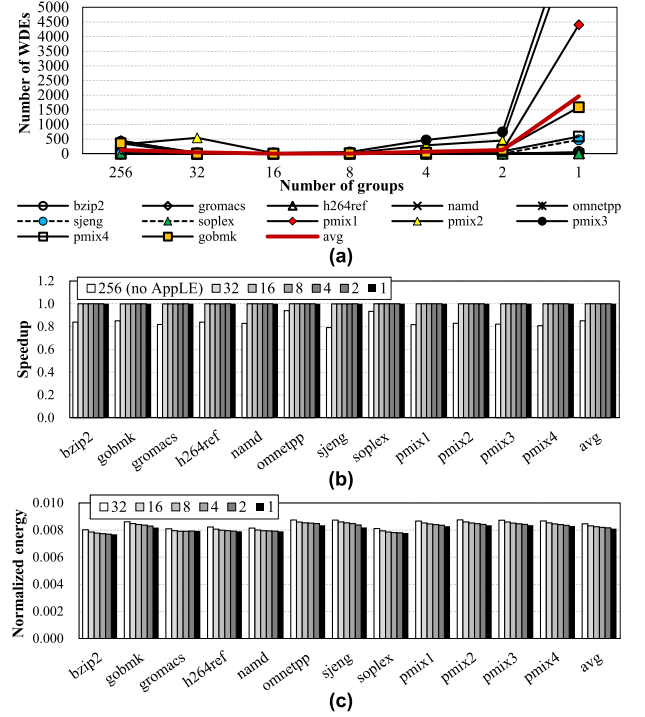


Fig. 11. Sensitivity to the number of groups in AppLE: (a) normalized WDE, (b) speedup, (c) normalized energy.

Fig. 11c shows the SRAM energy normalized to the case without AppLE. In general, the energy decreases as the number of read ports shrinks.

5.7 Comparison With Other Studies

From the sensitivity analysis above, the most cost-effective IMDB becomes IMDB(e256b8g8), which consists of 256 entries in the main table, 8 entries in the barrier buffer, and a group size of 8. The group size of 4 is denoted as IMDB(e256b8g4). Five schemes are compared against IMDB: (1) *PARR*, (2) *FnW* [4], (3) *Lazy correction* [40], (4) *ADAM* [39], and (5) *SIWC* [16].

PARA (probabilistic adjacent row activation) is commonly used for mitigating rowhammers in DRAM devices [21]. Preventing occurrences of WDEs requires restoration (i.e., rewrite) rather than activation; hence, rewrite commands for adjacent row data should be randomly generated when a normal write command goes into the media controller. This study evaluates *PARRs* (probabilistic adjacent row restoration) with different probabilities (i.e., $p = 0.1 - 0.0001$). *FnW* inverts the data if more than half of the bits are changed [4]; it can minimize the number of bit flips in a PCM device. Since *FnW* is a device-level approach, it is applied to the proposed scheme. *Lazy correction* defers subsequent VnC by temporarily storing errors in an ECP chip [40]. Each entry of ECP records multiple errors of one PCM line. We assume that 10 pointers, which is the maximum number in [40], are handled in the ECP. *ADAM* aligns the compressed data in the device alternately to avoid data pattern that is vulnerable to WDEs [39]. *SIWC* sparsely caches write data in an SRAM [16]. In particular, *SIWC-size* indicates that the SRAM capacity is identical to that of IMDB, and *SIWC-entry* holds entries in an amount equal to that of IMDB.

TABLE 5
Performance of Different Mitigation Schemes

Schemes	WDEs	Speedup	Energy
PARR(p=0.1)	41.915	0.9718	1.09468
PARR(p=0.01)	4.5090	0.9971	1.00947
PARR(p=0.001)	0.2670	0.9997	1.00095
PARR(p=0.0001)	0.7532	0.9999	1.00009
Lazy correction	0.19→0	0.362782	2.177345
ADAM	0.5341	0.9807	1.1765
SIWC-size	0.7276	1.0417	0.9467
SIWC-entry	0.0885	1.0628	0.8951
IMDB(e256b8g4)	2.08E-3	0.9561	0.9937
IMDB(e256b8g8)	4.39E-4	0.9560	0.9941
IMDB(e256b8g4)+FnW	1.66E-3	0.9560	0.9975
IMDB(e256b8g8)+FnW	1.97E-4	0.9560	0.9977

5.7.1 Write Disturbance Errors

The second column in Table 5 reports normalized WDEs. PARR shows lower WDEs as the probability scales down, except for $p=0.0001$. Since rewrite commands might be unnecessary on the infrequently accessed row, excessive restoration with high probability may incur more WDEs (i.e., 41.915 on $p=0.1$). The lowest probability of 0.0001 in Table 5 also leads to more WDEs, because restoration on vulnerable cells is scarce. *Lazy correction* yields non-zero normalized WDE values for different ECPs; however, it is noteworthy that *lazy correction* shows *temporal* WDEs in run-time, which can finally be corrected with ECPs. *SIWC-entry* presents 87.84% lower WDEs than *SIWC-size* (i.e., 0.0885 versus 0.7276) because the mitigation performance strongly depends on the cache size. *ADAM* is effective only if the compression ratio exceeds 0.5; hence, *ADAM* shows inferior performance.

In contrast, IMDB(e256b8g8) reduces WDEs to 4.39E-4, which is 1218 \times and 202 \times fewer WDEs compared to *ADAM* and *SIWC-entry*, respectively. It is noteworthy that these configurations show comparable WDE mitigation performance to the case where the main table consists of 2048 entries without barrier buffers. While a 2048-entry main table requires $108b \times 2048 \times 4\text{-bank} = 864\text{KB}$ of SRAM, the combinational approach yields fewer WDEs with a 16KB SRAM, which is four times smaller than *SIWC*. Furthermore, applying FnW to IMDB(e256b8g8) yields 2.2 \times fewer WDEs, due to a reduction in the number of bit flips.

5.7.2 Speedup

The third column in Table 5 presents the speedup compared to the baseline. PARR achieves similar performance with the baseline regardless of the restoration probability. *Lazy correction* shows the lowest speedup. This is because even though the VnC for corrupted data is deferred, at least four read operations strictly ordered by a write command are necessary. Although the proposed method rewrites two neighbors, these operations are performed in an on-demand fashion instead of incurring four read operations per write operation, as VnC does. Therefore, the proposed method can outperform *lazy correction*. The speed of *ADAM* degrades by about 2% due to encoding and decoding processes of FPC. For *SIWC-entry* and *-size*, slightly higher performance is achieved.

On the other hand, two configurations of the proposed method experience approximately 4% speed degradation on average. The waiting cycles for memory systems constitute 12% of execution time in the baseline, according to our evaluation. Consequently, the proposed method degrades the performance of the overall system only by 0.48%. IMDB requires 1-3 cycles for processing a write command on the critical path. The latency is determined by the hit/miss cases of the main table and the barrier buffer. If a write command hits on the main table, one referring cycle is spent on the main table. Furthermore, suppose this hit command triggers the rewrite operation. In that case, one more read cycle on the barrier buffer is required, because the hit entry must be promoted to the barrier buffer (i.e., contents swapping). Finally, the swapped contents are written to the main table and the barrier buffer, maximally resulting in three cycles. If a write command misses, AppLE must be performed for finding a replacing candidate. However, the latency of AppLE can be hidden within the write latency (i.e., IDLE state in Fig. 6c). In a memory system, all commands follow promised timing constraints (i.e., JEDEC DDR standards). Thus, the media controller must wait for the latency (i.e., 150 ns or 120 cycles) after issuing a write command to a bank. As a result, the negligible latency of IMDB leads to minor performance degradation.

5.7.3 Energy

The fourth column in Table 5 shows the normalized energy. PARRs show higher energy consumption than the baseline for all probabilities, because rewrite commands cause higher write energy consumption. However, the energy overhead is not notable, due to relatively low probabilities (i.e., ≤ 0.1) of PARRs. *Lazy correction* consumes $2.18\times$ higher energy than the baseline, because both execution time and the number of commands increase. Meanwhile, *SIWC-size* reduces 5% of energy compared to the baseline. This is because persistent workloads have relatively high locality due to cache line flush instructions, reducing write operations on frequently accessed addresses. Furthermore, the energy can be reduced by about 10.5% compared to the baseline with a larger number of entries, as declared by *SIWC-entry*; however, it should be noted that the WDE mitigation performance is not as excellent as it is with the proposed methods. Although IMDB(e256b8g8) presents 9% higher energy consumption compared to *SIWC-entry*, this outcome is still 0.59% smaller than the baseline. Even though the proposed scheme generates rewrite commands that may contribute to the energy consumption, the “tiny” barrier buffer reduces the write traffic with a 10.67% cache hit rate, leading to lower energy consumption than the baseline. In contrast, IMDB(e256b8g8) consumes 54.4% less energy than *lazy correction*.

6 DISCUSSION

Synergy With ECC Schemes. In general, error-correcting codes (ECC) are proactively being employed in memory products that have reliability-related problems. In our case, ECC logic is placed on the media controller for system expandability. To observe the system reliability, we evaluated failure-in-time (FIT), which is the number of corrupted bits in an hour [13], [28]. Commonly, Fig. 12 shows that FITs decrease when the correction capability of ECC enhances. In




Fig. 12. FITs when different ECC schemes are applied to (a) IMDB, (b) baseline, (c) IMDB(e256b8g8) for a 64Gb bank.

particular, Fig. 12a shows that 0-FIT can be achieved when ECC4 (i.e., 4-bit error correction) and ECC8 (i.e., 8-bit error correction) are applied to IMDB (e256b8g8) and IMDB (e256b8g4), respectively. A (552, 512)-BCH code that is capable of correcting 4 errors [42] only incurs 1.5ns of latency (i.e., < 1 cycle at 800MHz), according to the latency formula in [37]. Therefore, only a minuscule amount of latency is required when IMDB is assisted by ECC. Fig. 12b shows that ECC-16 is ineffective for WDEs. We observe that simultaneous bit flips occur in one data across all workloads, leading to simultaneous WDEs in multiple cells. Since ECC has no knowledge of such programming patterns, ECC is incapacitated by WDEs. In contrast, Fig. 12b shows that ECC256 yields 0-FIT. The correction capability of ECC256 corresponds to a (3584, 512)-BCH code. This code yields $611 \times$ larger area than that of (552, 512)-BCH code according to the area formula in [37]. Therefore, IMDB is necessary for obtaining a reliable memory system with a lower area burden on ECC.

Discussion of SRAM Capacity Against SIWC. Considering the capacity of SRAM for the proposed method and a write cache-based study (i.e., SIWC) in a four-bank PCM system, the latter requires $256 \times 64b \times 4\text{-bank} = 64\text{KB}$ of SRAM if 256 addresses are managed per bank. On the other hand, for the proposed method, the main table entry has $25b + 8b + 72b + 3b = 108b$, and the barrier buffer entry has $64b + 25b + 8b + 8b = 553b$ (see Fig. 3a). Therefore, the proposed method requires $256 \times 108b \approx 3.4\text{KB}$ of SRAM on the main table per PCM bank. We evaluate our system by configuring the main table as a fully associative SRAM, because a 3.4KB of fully associative SRAM cache does not burden resources. In addition, a fully associative can yield the best performance compared to fewer ways. The barrier buffer consumes $8 \times 553b \approx 0.6\text{KB}$ of SRAM per PCM bank (see Section 5.5). Consequently, $(3.4\text{KB} + 0.6\text{KB}) \times 4\text{-bank} = 16\text{KB}$ of SRAM translates to 2KB per 1GB of PCM. If 256 addresses are managed, the proposed method consumes $4 \times$ smaller SRAM area than SIWC, and the gap enlarges as the number of managed addresses grows. Besides the SRAM capacity, introducing SRAM as a data region requires considering the hold-up

time constraint of supercapacitors. In particular, SIWC only holds dirty data; hence, flushing 256 volatile data requires $150\text{ns} \times 256 \text{ flushes} / 100\mu\text{s} = 38.4\%$ of flush time at most (i.e., all row buffer miss commands on a single bank), where the value of $100\mu\text{s}$ comes from [15]. In contrast, flushing data in the barrier buffer only requires $150\text{ns} \times 8 \text{ flushes} / 100\mu\text{s} = 1.2\%$. In conclusion, IMDB mitigates more WDEs without expanding supercapacitors.

Discussion of Area Overhead. IMDB mainly consists of a main table, a barrier buffer, control logic for AppLE, and control logic for integrated counter blocks. First, the main table and the barrier buffer require 16 KB of SRAM, which translates to 768K transistors considering 6T SRAM. Second, the control logic for AppLE requires a 9-bit comparator. The comparator consists of one AND-gate, one NOR-gate, and 26 AND-gates with one bubbled input [27], each requiring 6, 4, and 10 transistors, respectively. Thus, the 9-bit comparator consists of a total of 270 transistors ($= 6 + 4 + 26 \times 10$). Lastly, the control logic for integrated counter blocks in Fig. 3b consists of 5.5M transistors according to synthesis results from Synopsys Design Compiler. Consequently, IMDB consists of 6.268M transistors in total. We find that a representative DRAM controller in [3] requires approximately 3.7B transistors (i.e., 1.8 mm^2 at 22 nm). Therefore, IMDB incurs 0.17% of area overhead with respect to the representative DRAM controller. It is noteworthy that the PCM controller area is not disclosed; however, the higher complexity of the PCM controller than that of DRAM explicitly proves that IMDB occupies a small amount of area.

Scalability of the Proposed Scheme. We evaluate WDEs for a larger bank density compared to a 2GB (i.e., 16Gb) bank that is adopted in Section 5 to observe the scalability of IMDB. The normalized WDE of a 64Gb bank with IMDB (e256b8g8) is $9.27\text{E-}3$, $20 \times$ higher than a 16Gb bank (i.e., $4.38\text{E-}4$). This result indicates that the currently proposed size has less effect on a larger density, because an IMDB plane should manage more addresses. We can address such a scalability issue in two ways. First, enlarging the number of entries in the main table to 512 achieves $4.54\text{E-}4$ WDEs, which is again similar to the result of the original 16Gb bank with IMDB(e256b8g8) (i.e., $4.38\text{E-}4$). Second, stronger ECC can be applied to mitigate WDE in a larger density bank. Fig. 12c presents FITs when various ECC schemes are applied to the 64 Gb bank PCM, which is supported by IMDB(e256b8g8). For achieving 0-FIT in a 64Gb PCM, ECC16 is necessary rather than ECC4, which is effective in the 16 Gb bank PCM (see Fig. 12a).

7 RELATED WORK

VnC-Based Schemes. VnC is the most solid method capable of preventing WDEs [38], [40]; it triggers two pre-write read operations and two post-write read operations, before and after a write operation, respectively. These four read commands are strictly ordered by one write command, incurring significant performance overhead. In [40], *lazy correction* temporarily stores the locations of disturbed cells in an error-correction pointer (ECP) chip, deferring the subsequent VnC as late as possible until the ECP becomes full. However, cells in the ECP must be well insulated to guarantee no errors. Also, it is necessary to execute at least four read operations for the initial write command.

Encoding-Based Schemes. Data encoding can reduce WDE-vulnerable patterns [10], [11], [12], [18], [36], [39]. In [18], DIN proposes a codebook that encodes contiguous 0s in a compressed pattern to eliminate patterns vulnerable to WDEs. However, this approach must fall back on the VnC method if the length of the encoded data exceeds the length of the cache line. In [10], MinWD encodes write data into multiple candidates with special shift operations and selects the least aggressive form from all candidates. However, this method requires additional bits as an indicator of the shift operation. In [39], ADAM compresses a cache line and aligns the line to the right and left alternately; hence, the number of valid bits on adjacent rows is reduced. However, encoding schemes strongly depend on the data patterns of the applications. WLC [36] is a compression scheme for reducing energy; it compresses few MSBs of each 64-bit of a cache line, increasing “the number lines” to be compressed. However, compared with the compression ratio of 40% in ADAM, the compression ratio of WLC is bounded to $9 \times 8 / 512 = 14.1\%$ if 9 MSBs in each 64-bit can be compressed. Thus, WLC is less effective than ADAM.

Cache-Based Scheme. Storing frequently updated data in volatile caches can enhance the system reliability. In [16], SIWC leverages a write cache that inserts data probabilistically and absorbs bit flips. Because WDE-vulnerable data would be stored in the write cache, the victims of WDEs become safe. However, this method introduces several mega-bytes of volatile memory to obtain a high hit ratio, and the supercapacitor for data flushes must be expanded. Furthermore, SIWC reports the number of operations that may incur WDEs (i.e., WDE limitation number), but this information is not utilized for WDE mitigation.

8 CONCLUSION

WDE is a severe reliability problem that hinders the manufacturing of PCMs. This study proposes a table-based approach, IMDB, to restore cells on demand within a module. The newly proposed replacement policy yields higher reliability than the LRU and fully randomized replacement policies. Subsequently, AppLE enables an efficient implementation of the replacement policy. The small barrier buffer absorbs bit flips, offloading the burden onto the supercapacitor. Consequently, rigorous sensitivity analyses concerning design parameters are conducted to obtain a cost-effective architecture. The evaluation results show that the proposed method significantly reduces WDEs compared to the outcomes of earlier studies while maintaining speed and energy consumption levels that approximate those of the baseline.

ACKNOWLEDGMENTS

We thank the anonymous reviewers to substantially improve the paper. Also, special appreciations go to Jiwoong Choi, Boyeal Kim, and Taehyun Kim for their feedback.

REFERENCES

- [1] ARM Cortex A-15 (Samsung Exynos 5250), 2012. [Online]. Available: <https://www.7-cpu.com/cpu/Cortex-A15.html>
- [2] N. Binkert et al., “The gem5 simulator,” *SIGARCH Comput. Archit. News*, vol. 39, no. 2, pp. 1–7, Aug. 2011.
- [3] M. N. Bojnordi and E. Ipek, “PARDIS: A programmable memory controller for the DDRx interfacing standards,” in *Proc. 39th Annu. Int. Symp. Comput. Archit.*, 2012, pp. 13–24.
- [4] S. Cho and H. Lee, “Flip-N-Write: A simple deterministic technique to improve PRAM write performance, energy and endurance,” in *Proc. 42nd Annu. Int. Symp. Microarchit.*, 2009, pp. 347–357.
- [5] J. Choi, J. Jang, and L. Kim, “DC-PCM: Mitigating PCM write disturbance with low performance overhead by using detection cells,” *IEEE Trans. Comput.*, vol. 68, no. 12, pp. 1741–1754, Dec. 2019.
- [6] J. Coburn et al., “NV-Heaps: Making persistent objects fast and safe with next-generation, non-volatile memories,” in *Proc. 16th Int. Conf. Architect. Support Program. Lang. Oper. Syst.*, 2011, pp. 105–118.
- [7] X. Dong, C. Xu, Y. Xie, and N. P. Jouppi, “NVSIM: A circuit-level performance, energy, and area model for emerging nonvolatile memory,” *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 31, no. 7, pp. 994–1007, Jul. 2012.
- [8] J. L. Henning, “SPEC CPU2006 benchmark descriptions,” *SIGARCH Comput. Archit. News*, vol. 34, no. 4, pp. 1–17, Sep. 2006.
- [9] Hewlett Packard Lab. Cacti 6.5, 2008. [Online]. Available: <https://www.hpl.hp.com/research/cacti/>
- [10] M. Imran, T. Kwon, and J.-S. Yang, “Effective write disturbance mitigation encoding scheme for high-density PCM,” in *Proc. Des. Autom. Test Eur. Conf. Exhib.*, 2020, pp. 1490–1495.
- [11] M. Imran, T. Kwon, N. A. Toubia, and J.-S. Yang, “CEnT: An efficient architecture to eliminate intra-array write disturbance in PCM,” *IEEE Trans. Comput.*, vol. 71, no. 5, pp. 992–1007, May 2022.
- [12] M. Imran, T. Kwon, and J.-S. Yang, “ADAPT: A write disturbance aware programming technique for scaled phase change memory,” *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 41, no. 4, pp. 950–963, Apr. 2022.
- [13] T. Instruments, “Reliability terminology,” 2021. [Online]. Available: <https://www.ti.com/support-quality/reliability/reliability-terminology.html>
- [14] Intel, “Intel memory drive technology set up and configuration guide,” 2019. [Online]. Available: <https://www.intel.com/content/dam/support/us/en/documents/memory-and-storage/intel-mdt-setup-guide.pdf>
- [15] J. Izraelevitz et al., “Basic performance measurements of the Intel Optane DC persistent memory module,” 2019, *arXiv: 1903.05714*.
- [16] J. Jang, W. Shin, J. Choi, Y. Kim, and L.-S. Kim, “Sparse-insertion write cache to mitigate write disturbance errors in phase change memory,” *IEEE Trans. Comput.*, vol. 68, no. 5, pp. 752–764, May 2019.
- [17] J. Jeong et al., “Efficient hardware-assisted logging with asynchronous and direct-update for persistent memory,” in *Proc. 51st Int. Symp. Microarchit.*, 2018, pp. 520–532.
- [18] L. Jiang, Y. Zhang, and J. Yang, “Mitigating write disturbance in super-dense phase change memories,” in *Proc. 44th Annu. IEEE/IFIP Int. Conf. Dependable Syst. Netw.*, 2014, pp. 216–227.
- [19] A. Joshi, V. Nagarajan, S. Viglas, and M. Cintra, “ATOM: Atomic durability in non-volatile memory through hardware logging,” in *Proc. 23rd Int. Symp. High Perform. Comput. Archit.*, 2017, pp. 361–372.
- [20] M. Kim, H. Lee, H. Kim, and H.-J. Lee, “WL-WD: Wear-leveling solution to mitigate write disturbance errors for phase-change memory,” *IEEE Access*, vol. 10, pp. 11 420–11 431, 2022.
- [21] Y. Kim et al., “Flipping bits in memory without accessing them: An experimental study of DRAM disturbance errors,” in *Proc. 41st Int. Symp. Comput. Archit.*, 2014, pp. 361–372.
- [22] Y. Kim, S. Yoo, and S. Lee, “Write performance improvement by hiding R drift latency in phase-change RAM,” in *Proc. DAC Des. Automat. Conf.*, 2012, pp. 897–906.
- [23] B. C. Lee et al., “Architecting phase change memory as a scalable DRAM alternative,” in *Proc. 36th Annu. Int. Symp. Comput. Archit.*, 2009, pp. 2–13.
- [24] H. Lee, M. Kim, H. Kim, H. Kim, and H. Lee, “Integration and boost of a read-modify-write module in phase change memory system,” *IEEE Trans. Comput.*, vol. 68, no. 12, pp. 1772–1784, Dec. 2019.
- [25] S. H. Lee, “Method of driving phase change memory device capable of reducing heat disturbance,” U.S. Patent 0 204 664, 2014.
- [26] Mentor-Graphics, “DDR4 and LPDDR4 - broad design verification and challenges,” 2013. [Online]. Available: <https://www.mentor.com/pcb/multimedia/player/ddr4-and-lpddr4-board-design-verification-and-challenges-356bbc16-6195-4d78-ba85-5496362bec44>
- [27] MICREL, 2007. [Online]. Available: <https://www.mouser.com/datasheet/2/268/sy100s366-778966.pdf>

- [28] S. S. Mukherjee et al., "The soft error problem: An architectural perspective," in *Proc. 11th Int. Symp. High Perform. Comput. Archit.*, 2005, pp. 243–247.
- [29] P. J. Nair et al., "Reducing read latency of phase change memory via early read and turbo read," in *Proc. 21st Int. Symp. High-Perform. Comput. Archit.*, 2015, pp. 309–319.
- [30] I. B. Peng et al., "System evaluation of the Intel optane byte-addressable NVM," in *Proc. Int. Symp. Memory Syst.*, 2019, pp. 304–315.
- [31] M. Poremba, T. Zhang, and Y. Xie, "NVMain 2.0: A user-friendly memory simulator to model (non-)volatile memory systems," *IEEE Comput. Archit. Lett.*, vol. 14, no. 2, pp. 140–143, Jul.–Dec. 2015.
- [32] M. K. Qureshi, A. Sezenc, L. A. Lastras, and M. M. Franceschini, "Practical and secure PCM systems by online detection of malicious write streams," in *Proc. 17th Int. Symp. High-Perform. Comput. Archit.*, 2011, pp. 478–489.
- [33] M. K. Qureshi et al., "Enhancing lifetime and security of PCM-based main memory with start-gap wear leveling," in *Proc. 42nd Annu. Int. Symp. Microarchit.*, 2009, pp. 14–23.
- [34] U. Russo, D. Ielmini, A. Redaelli, and A. L. Lacaita, "Modeling of programming and read performance in phase-change memories—Part II: Program disturb and mixed-scaling approach," *IEEE Trans. Electron Devices*, vol. 55, no. 2, pp. 515–522, Feb. 2008.
- [35] N. H. Seong, D. H. Woo, and H.-H. Lee, "Security refresh: Protecting phase-change memory against malicious wear out," *IEEE Micro*, vol. 31, no. 1, pp. 119–127, Jan./Feb. 2011.
- [36] S. Seyedzadeh et al., "Enabling fine-grain restricted coset coding through word-level compression for PCM," in *Proc. 24th Int. Symp. High Perform. Comput. Archit.*, 2018, pp. 350–361.
- [37] D. Strukov, "The area and latency tradeoffs of binary bit-parallel BCH decoders for prospective nanoelectronic memories," in *Proc. 40th Asilomar Conf. Signals Syst. Comput.*, 2006, pp. 1183–1187.
- [38] H. Sun et al., "Design techniques to improve the device write margin for MRAM-based cache memory," in *Proc. 21st Ed. Great Lakes Symp. VLSI*, 2011, pp. 97–102.
- [39] S. Swami and K. Mohanram, "ADAM: Architecture for write disturbance mitigation in scaled phase change memory," in *Proc. Des. Autom. Test Eur. Conf. Exhib.*, 2018, pp. 1235–1240.
- [40] R. Wang et al., "SD-PCM: Constructing reliable super dense phase change memory under write disturbance," in *Proc. 20th Int. Conf. Architect. Support Program. Lang. Operating Syst.*, 2015, pp. 19–31.
- [41] M. Weiland et al., "An early evaluation of Intel's optane DC persistent memory module and its impact on high-performance scientific applications," in *Proc. Int. Conf. High Perform. Comput. Netw. Storage Anal.*, 2019, Art. no. 76.
- [42] C. Yang et al., "Improving reliability of non-volatile memory technologies through circuit level techniques and error control coding," *EURASIP J. Adv. Signal Process.*, vol. 212, no. 1, pp. 1–24, 2012.
- [43] J. Yang et al., "An empirical guide to the behavior and use of scalable persistent memory," in *Proc. 18th USENIX Conf. File Storage Technol.*, 2020, pp. 169–182.
- [44] H. Yu and Y. Du, "Increasing endurance and security of phase-change memory with multi-way wear-leveling," *IEEE Trans. Comput.*, vol. 63, no. 5, pp. 1157–1168, May 2014.



Hyokeun Lee (Member, IEEE) received the BS and PhD degrees in electrical and computer engineering (ECE) from Seoul National University, Seoul, South Korea, in 2016 and 2021, respectively. He is currently working as a postdoctoral researcher with Inter-University Semiconductor Center (ISRC), Seoul National University. His current research interests include non-volatile memory controller design, architecture simulation modeling, and computer architecture.



Seungyong Lee received the BS degree in electrical and computer engineering (ECE) from Seoul National University, Seoul, South Korea, in 2018. He is currently working toward the integrated MS and PhD degrees in electrical and computer engineering with Seoul National University. His current research interests include processing-in-memory, memory controller design, and computer architecture.



Byeongki Song received the BS and MS degrees in electrical and computer engineering from Seoul National University, Seoul, South Korea, in 2019, and 2021 respectively. From 2021, he works as an engineer with Samsung Electronics, Hwaseong, Gyeonggi-do, South Korea. His current research interests include low-power SoC design for multimedia applications, computer architecture, and deep learning.



Moonsoo Kim received the BS and PhD degrees in electrical and computer engineering from Seoul National University, Seoul, South Korea, in 2014 and 2020, respectively. In 2020, he joined SoC design team of Samsung Electronics, Hwasung, South Korea, where he is currently working as a staff engineer. His research interests include the areas of cache/memory architecture, and SoC design.



Seokbo Shim received the BS degree in electrical engineering (EE) from Korea University, in 2003, and the MS degree in electrical and computer engineering (ECE) from Seoul National University, Seoul, South Korea, in 2021. In 2004, he worked with LG Electronics. From 2005 to 2019, he worked as a principal engineer with SK Hynix Inc., Icheon, South Korea. He is currently working as a principal engineer with SK Hynix Inc. His current research interests include DDR4/DDR5 memory design architecture.



Hyuk-Jae Lee (Member, IEEE) received the BS and MS degrees in electrical engineering (EE) from Seoul National University, Seoul, South Korea, in 1987 and 1989, respectively, and the PhD degree in electrical and computer engineering (ECE) from Purdue University, West Lafayette, Indiana, in 1996. From 1998 to 2001, he was a senior component design engineer with the Server and Workstation Chipset Division, Intel Corporation, Hillsboro, Oregon. From 1996 to 1998, he was a faculty member with the Department of Computer Science, Louisiana Tech University, Ruston, Louisiana. In 2001, he joined the School of EECS, Seoul National University, where he is a professor. He is the founder of Mamurian Design, Inc., Seoul, a fabless SoC design house for multimedia applications. His current research interests include computer architecture and SoC for multimedia applications.



Hyun Kim (Member, IEEE) received the BS, MS, and PhD degrees in electrical engineering and computer science (EECS) from Seoul National University, South Korea, in 2009, 2011, and 2015, respectively. From 2015 to 2018, he was a BK assistant professor with the BK21 Creative Research Engineer Development for IT, Seoul National University. In 2018, he joined the Department of Electrical and Information Engineering, Seoul National University of Science and Technology, where he is an assistant professor. His current research interests include algorithm, computer architecture, memory, and SoC design for low-complexity multimedia applications, and deep neural networks.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.