Microprocessors and Microsystems 57 (2018) 65–75

Contents lists available at ScienceDirect

Microprocessors and Microsystems

journal homepage: www.elsevier.com/locate/micpro

Adaptive-Classification CLOCK: Page replacement policy based on read/write access pattern for hybrid DRAM and PCM main memory

Sungho Kim[a], Sang-Ho Hwang[b], Jong Wook Kwak[a,*]

[a] Department of Computer Engineering, Yeungnam University, Gyeongsan, Republic of Korea
[b] Daegu Gyeongbuk Institute of Science and Technology, Daegu, Republic of Korea

ARTICLE INFO

Keywords:
Hybrid main memory
Phase change memory
Non-volatile memory
AC-CLOCK
CLOCK
Read/write access pattern

ABSTRACT

Phase change memory (PCM) has emerged as a new main memory in embedded systems because of its attractive characteristics, which include a lower static power, higher density, byte-addressability, and non-volatility. Although PCM has various advantages, it has a longer write latency than dynamic random access memory (DRAM) and limits the write endurance in each cell. Recently, a hybrid main memory was designed that combined DRAM and PCM. In particular, the page replacement policy in this hybrid main memory has been one of the major areas of study. The previous page replacement policies had three problems when used in embedded systems: the allocation policy, read/write access pattern identification, and migration policy. This paper proposes the adaptive-classification CLOCK (AC-CLOCK) page replacement policy for the hybrid main memory of an embedded system. AC-CLOCK measures the read/write access patterns of application programs and reduces the number of write operations in PCM by tracing these patterns. In addition, it reduces the number of unnecessary page migrations from PCM to DRAM by tracking the detailed access patterns such as write intensive pages or read intensive pages. In our experiments, AC-CLOCK reduced the number of write operations in PCM by 57.87% on average and outperformed the hit ratio of write operations in the DRAM by 9.47% on average, compared with CLOCK, CLOCK-Pro, and CLOCK-DWF. Moreover, AC-CLOCK reduced the energy delay product (EDP) by 23.25% on average, compared with the previous techniques.

1. Introduction

Over the past several years, there has been an explosive increase in the demand for embedded systems for use in the Internet of Things (IoT), smart phones, and wearable devices [1–5]. As a result, applications in these embedded systems have required more hardware resources, especially increases in the capacity and speed of main memories [6]. Dynamic random access memory (DRAM) has generally been used as the main memory in embedded systems [7–9]. However, although it has the advantage of a shorter read/write latency, it consumes more static power (e.g., for refresh operations) to maintain the resources of the applications residing in the main memory. Previous studies found that the DRAM main memory accounted for 30–48% of all the energy consumed in embedded systems [10].

To reduce the energy consumption of the main memory in embedded systems, phase change memory (PCM) has emerged as a new main memory in academia and industry [11–14]. While it has attractive characteristics such as a lower static power, higher density, byte-addressability, and non-volatility, it has a longer write latency than DRAM and limits the write endurance in each cell to approximately 10^6–10^8 [14]. Moreover, a PCM consumes approximately 10 times (1.0 nj/bit) more write energy than a DRAM [11]. Although previous studies have made technological advances such as reducing the write latency in PCM by modifying the hardware designs and have improved the lifetime of PCM using additional hardware logic, PCM cannot yet totally replace the traditional DRAM [13–20].

Recently, to exploit the advantages of PCM, hybrid main memories with a DRAM and PCM combination have been designed [1–3,14,21]. In this hybrid main memory, the paging system is an important part of the virtual memory implementation. Paging is a technique that allows programs that exceed the size of the available physical main memory to run correctly. Therefore, many researchers have mostly focused on studying the paging system and a page replacement policy to manage the faulted pages of the paging technique [22–26]. Hybrid main memories are classified using two organizational methods: a hierarchical organization and parallel organization. Fig. 1 shows these two organizational methods for a hybrid main memory [14].

The hierarchical organization uses the DRAM cache as the upper

* Corresponding author.
E-mail address: kwak@yu.ac.kr (J.W. Kwak).

https://doi.org/10.1016/j.micpro.2018.01.003
Received 24 August 2017; Received in revised form 15 December 2017; Accepted 19 January 2018
Available online 31 January 2018

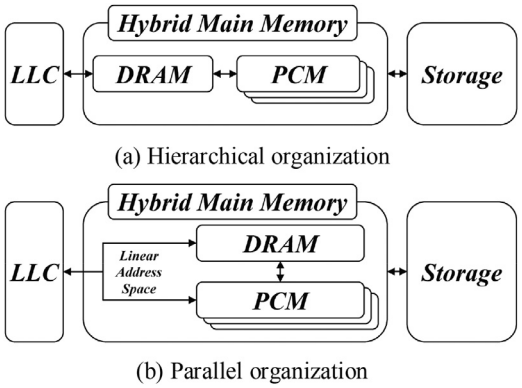

Fig. 1. Organizations of hybrid main memory [14].

PCM main memory. It can still utilize an intact page replacement policy based on the DRAM such as LRU or CLOCK [22–26]. However, the hierarchical organization has an extra hardware cost for this additional layer, which is the DRAM cache. Moreover, a data reliability technique is still required between the DRAM cache and PCM main memory [13–14]. By organizing DRAM and PCM on the same layer, the parallel organization can use a fully associative replacement because a page is directly managed by a kernel, unlike the hierarchical organization, which does not impose an extra hardware cost. However, it cannot directly utilize page replacement policies based on the DRAM [27–28].

Because of advantages such as no additional hardware cost and the direct management by a kernel with a slight software modification in the parallel organization of the hybrid main memory, many researchers have tried to reduce the number of write operations in PCM using a page replacement policy based on the parallel organization. However, their techniques have insensitive characteristics with a change in the read/write access patterns of application programs [29–33]. The three major problems of their techniques are the allocation policy, read/write access pattern identification, and migration policy.

In this paper, we propose a new page replacement policy for mitigating the weaknesses of the previous techniques based on the parallel organization, which is called adaptive-classification CLOCK (AC-CLOCK). AC-CLOCK adopts new policies for classifying the detailed read/write access patterns in the residing pages, which solves the three problems of the previous techniques. The contributions of this paper are summarized as the following specific characteristics.

- AC-CLOCK has a high hit ratio of write operations in the DRAM, resulting in a reduction of the PCM writes, which positively affects the lifetime extension of the hybrid memory. By precisely tracking the read/write access patterns of all the pages, AC-CLOCK allocates the pages that have a high frequency of write operations to the DRAM.
- AC-CLOCK reduced the number of unnecessary page migrations from PCM to DRAM and thus reduces the execution latency of the hybrid memory system. This is the result of using two policies: the allocation policy and migration policy of AC-CLOCK. The allocation policy allocates pages to PCM by selecting pages with a low migration possibility from PCM to DRAM. In the migration policy, the pages allocated to PCM migrate to DRAM using a precise read/write access pattern identification.
- AC-CLOCK provides an energy efficient hybrid memory system, in terms of the energy delay product (EDP). In the hybrid main memory, the energy consumption is related to the number of write operations in PCM. By reducing the frequency of write operations in PCM, AC-CLOCK reduces the energy consumption of the hybrid main memory.

The remainder of this paper is organized as follows. Section 2 describes the background and related works for the hybrid main memory. In Section 3, we propose the AC-CLOCK page replacement policy. The performance evaluation of AC-CLOCK is shown in Section 4. Finally, Section 5 concludes this paper.

2. Background and related work

In this section, we explore the background of the PCM characteristics and describe related work on page replacement policies for the hybrid main memory.

Recently, PCM has been actively studied to replace the DRAM in embedded systems because of its attractive characteristics, which include a lower static power, higher density, byte-addressability, and non-volatility [1–5]. It has two states as a result of using a phase change material (e.g., $GeSeTe_2$ or Ag-In-Sb-Te) such as a chalcogenide glass: a crystalline state and an amorphous state [11]. In PCM, the phase-change material is crystallized by heating for a long period above a crystalline temperature, which is called the crystalline state (low resistance). In contrast, it is melt-quenched to achieve the amorphous temperature, which is called the amorphous state (high resistance). The crystalline state and amorphous state represent SET (value 1) and RESET (value 0), respectively. Fig. 2(a) shows the structure of a PCM cell [11]. The decision processes for the two states in the PCM cell are shown in Fig. 2(b) [11].

Although PCM has many advantages such as a lower static power, higher density, and byte addressability, it has a longer write latency than a DRAM and limits the write endurance in each cell to approximately 10^6–10^8 [14]. Moreover, PCM consumes approximately 10 times more write energy than a DRAM. To overcome these problems, previous researchers have tried to solve the weaknesses of PCM by reducing its write latency through modifying the hardware designs and improving its lifetime using additional hardware logic. Unfortunately, although these techniques have provided technological improvements, PCM cannot completely replace the traditional DRAM.

Recently, many researchers have proposed a new architecture that combines DRAM and PCM, which is called a hybrid main memory, to exploit the advantages of PCM [1–3,14,21]. They have also proposed page replacement policies to efficiently manage this hybrid main memory. Jiang et al. proposed a page replacement policy based on CLOCK for a DRAM, which is called CLOCK-pro [26]. They tried to solve the scan and loop problems in application programs, in the case of using page replacement policies such as the traditional LRU and CLOCK. CLOCK-pro was designed with the objective of identifying hot pages and cold pages in the main memory. To classify pages, they defined hot pages and cold pages by the frequencies that they were recently accessed. The cold pages were referenced only once recently. In contrast, the hot pages were referenced at least twice recently. Moreover, to predict future access patterns, CLOCK-pro used the history access list information of recently evicted pages. When there was a history access of a replacing page in the history access list, their technique allocated a hot page to the replacing page. Otherwise, it allocated a cold page to the replacing page. However, because CLOCK-pro was designed as the page replacement policy for DRAM, it did not consider

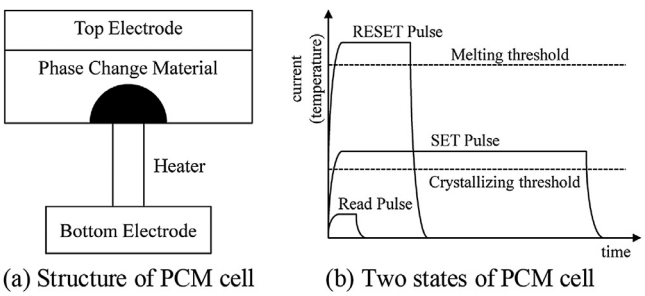

Fig. 2. Structure and two states of PCM cell [11].