

**TOBB UNIVERSITY OF ECONOMICS AND TECHNOLOGY**  
**INSTITUTE OF NATURAL AND APPLIED SCIENCES**

**REDUCING DRAM ACCESS LATENCY BY EXPLOITING DRAM LEAKAGE  
CHARACTERISTICS AND COMMON ACCESS PATTERNS**



**MASTERS THESIS**

**Hasan HASSAN**

**Department of Computer Engineering**

**Supervisor: Assoc. Prof. Oguz ERGIN**

**AUGUST 2016**



Approval of the Institute of Natural and Applied Sciences

.....  
**Prof. Osman EROGUL**  
Director

I certify that this thesis satisfies all the requirements as a thesis for the degree of Master of Science.

.....  
**Assoc. Prof. Oguz ERGIN**  
Deputy Head of Department

This thesis entitled "**REDUCING DRAM ACCESS LATENCY BY EXPLOITING DRAM LEAKAGE CHARACTERISTICS AND COMMON ACCESS PATTERNS**" has been prepared and submitted in partial fulfillment of the requirements for the degree of Master of Science in Computer Engineering by **Hasan HASSAN**, who is a graduate student at TOBB University of Economics and Technology, Institute of Natural and Applied Sciences with student number 131111040. The thesis has been examined in **AUGUST 11, 2016** by the thesis committee below and is recommended for approval and acceptance.

**Supervisor:**

**Assoc. Prof. Oguz ERGIN** .....  
TOBB University of Economics and Technology

**Committee Members:**

**Prof. Mehmet Onder EFE (Chair)** .....  
Hacettepe University

**Assoc. Prof. Ali BOZBEY** .....  
TOBB University of Economics and Technology



## **TEZ BİLDİRİMİ**

Tez içindeki bütün bilgilerin etik davranış ve akademik kurallar çerçevesinde elde edilerek sunulduğunu, alıntı yapılan kaynaklara eksiksiz atıf yapıldığını, referansların tam olarak belirtildiğini ve ayrıca bu tezin TOBB ETÜ Fen Bilimleri Enstitüsü tez yazım kurallarına uygun olarak hazırlandığını bildiririm.

## **DECLARATION**

I hereby declare that all the information provided in this thesis has been obtained with rules of ethical and academic conduct and has been written in accordance with thesis format regulations. I also declare that, as required by these rules and conduct, I have fully cited and referenced all material and results that are not original to this work.

Hasan HASSAN



## ABSTRACT

Master of Science

### REDUCING DRAM ACCESS LATENCY BY EXPLOITING DRAM LEAKAGE CHARACTERISTICS AND COMMON ACCESS PATTERNS

Hasan HASSAN

TOBB University of Economics and Technology  
Institute of Natural and Applied Sciences  
Department of Computer Engineering

Supervisor: Assoc. Prof. Oguz ERGIN

Date: AUGUST 2016

DRAM-based memory is a critical factor that creates a bottleneck on the system performance since the processor speed largely outperforms the DRAM latency. In this thesis, we develop a low-cost mechanism, called *ChargeCache*, which enables faster access to recently-accessed rows in DRAM, with no modifications to DRAM chips. Our mechanism is based on the key observation that a recently-accessed row has more charge and thus the following access to the same row can be performed faster. To exploit this observation, we propose to track the addresses of recently-accessed rows in a table in the memory controller. If a later DRAM request hits in that table, the memory controller uses lower timing parameters, leading to reduced DRAM latency. Row addresses are removed from the table after a specified duration to ensure rows that have leaked too much charge are not accessed with lower latency. We evaluate ChargeCache on a wide variety of workloads and show that it provides significant performance and energy benefits for both single-core and multi-core systems.

**Keywords:** Dynamic Random Access Memory (DRAM), Memory systems.





## ÖZET

Yüksek Lisans Tezi

### DRAM SIZMA KARAKTERİSTİKLERİ VE OLAĞAN ERİŞİM ÖRÜNTÜSÜNDEN FAYDALANARAK DRAM ERİŞİM GECİKMESİNİN AZALTILMASI

Hasan HASSAN

TOBB Ekonomi ve Teknoloji Üniversitesi  
Fen Bilimleri Enstitüsü  
Bilgisayar Mühendisliği Anabilim Dalı

Supervisor: Doç. Dr. Oğuz ERGİN

Tarih: AĞUSTOS 2016

DRAM tabanlı bellek, bilgisayar sisteminde darboğaz oluşturarak sistemin başarımı sınırlayan en önemli bileşendir. Bunun sebebi işlemcilerin hız bakımından DRAM'lerin çok önünde olmasıdır. Bu tezde, *ChargeCache* ismini verdiğimiz, DRAM'lerin erişim gecikmesini azaltan bir yöntem geliştirdik. Bu yöntem, piyasadaki DRAM yongalarının mimarisinde bir değişiklik gerektirmediği gibi, bellek denetimcisinde de düşük donanım maliyeti olan ek birimlere ihtiyaç duymaktadır. *ChargeCache*, yeni erişilmiş DRAM satırlarının kısa bir süre sonra tekrar erişileceği gözlemine dayanmaktadır. Yeni erişilmiş satırlardaki DRAM hücreleri yüksek miktarda yük içerdiğinden, bunlara hızlı bir şekilde erişilebilir. Bu gözlemden faydalanmak için yeni erişilen satırların adreslerini bellek denetimcisi içerisinde bir tabloda tutmayı öneriyoruz. Sonraki erişim isteklerinin bu tablodaki satırlara erişmek istemesi durumunda, bellek denetimcisi yük miktarı yüksek hücrelerin erişilmek üzere olduğunu bileceğinden, DRAM erişim değiştirgelerini ayarlayarak erişimin düşük gecikmeyle tamamlanmasını sağlayabilir. Belirli bir süre sonra tablodaki satır adresleri silinerek, zaman içerisinde çok fazla yük kaybedip hızlı erişilebilme özelliğini yitirmiş satırların bu tablodan çıkarılması sağlanır. Önerdiğimiz yöntemi hem tek çekirdekli hem de çok çekirdekli mimarilerde benzetim ortamında deneyerek, yöntemin başarımla ve enerji kullanımı açısından sistem üzerinde sağladığı iyileştirmeleri inceledik.

**Anahtar Kelimeler:** Devingen Rastgele Erişimli Bellek, Bellek sistemleri.



## **ACKNOWLEDGMENTS**

I would like to thank my advisor Oguz Ergin for supporting me in every aspect through my undergraduate and graduate education in TOBB University of Economics and Technology. I would not have succeed without the priceless knowledge I acquired by working with him. I would also like to thank Onur Mutlu and SAFARI for for all the feedback and comments which greatly enhanced my research, KASIRGA for creating a stimulating working environment, and TOBB University of Economics and Technology for funding me during my education.



## TABLE OF CONTENTS

	<u>Page</u>
<b>ABSTRACT</b> . . . . .	vii
<b>ÖZET</b> . . . . .	ix
<b>ACKNOWLEDGMENTS</b> . . . . .	xi
<b>TABLE OF CONTENTS</b> . . . . .	xiii
<b>LIST OF FIGURES</b> . . . . .	xv
<b>LIST OF TABLES</b> . . . . .	xvii
<b>ABBREVIATIONS</b> . . . . .	xix
<b>1. INTRODUCTION</b> . . . . .	1
<b>2. BACKGROUND ON MAIN MEMORY</b> . . . . .	5
2.1 DRAM Organization . . . . .	6
2.1.1 Channel . . . . .	6
2.1.2 Rank . . . . .	7
2.1.3 Bank . . . . .	7
2.1.4 Subarray and row . . . . .	9
2.1.5 Cell . . . . .	9
2.2 DRAM Standards . . . . .	10
2.2.1 Double data rate type 3 (DDR3) . . . . .	10
2.3 DDR3 Operation . . . . .	10
2.4 Memory Controller . . . . .	12
<b>3. MOTIVATION</b> . . . . .	15
<b>4. CHARGE CACHE</b> . . . . .	19
4.1 High-level Overview . . . . .	19
4.2 Detailed Design . . . . .	19
4.2.1 Inserting rows into HCRAC . . . . .	20
4.2.2 Employing lowered DRAM timing constraints . . . . .	21
4.2.3 Invalidating stale rows from HCRAC . . . . .	21
4.3 Reduction in DRAM Timing Parameters . . . . .	21
<b>5. METHODOLOGY</b> . . . . .	23
<b>6. EVALUATION</b> . . . . .	25
6.1 Impact on Performance . . . . .	26
6.2 Impact on DRAM Energy . . . . .	26
6.3 Area and Power Consumption Overhead . . . . .	28
6.4 Sensitivity Studies . . . . .	29
6.4.1 ChargeCache capacity . . . . .	29
6.4.2 Caching duration . . . . .	29
<b>7. DISCUSSION</b> . . . . .	33
7.1 Temperature Independence . . . . .	33
7.2 Applicability to Other DRAM Standards . . . . .	33
<b>8. RELATED WORK</b> . . . . .	35
<b>9. CONCLUSION</b> . . . . .	37

<b>REFERENCES . . . . .</b>	<b>38</b>
<b>CURRICULUM VITAE . . . . .</b>	<b>47</b>



## LIST OF FIGURES

	<u>Page</u>
Figure2.1: Memory system of a modern computer. . . . .	5
Figure2.2: Layers of the DRAM hierarchy. . . . .	6
Figure2.3: View of a system with two channels. . . . .	7
Figure2.4: A channel which has two ranks that share data, command, and address buses. . . . .	8
Figure2.5: The internal structure of a rank which has 8 banks. . . . .	8
Figure2.6: The internal structure of a bank. . . . .	9
Figure2.7: Commands that are used to read data from DRAM and the timing parameters associated with them . . . . .	11
Figure2.8: Overview of a typical memory controller. . . . .	13
Figure3.1: Fraction of row activations that happen 8ms after precharge (8ms-RLTL) or refresh of the row ((a) Single-core workloads, (b) Eight-core workloads). . . . .	16
Figure3.2: RLTL for various time intervals ((a) Single-core workloads, (b) Eight-core workloads). . . . .	18
Figure4.1: Components of the ChargeCache Mechanism . . . . .	20
Figure4.2: Effect of initial cell charge on bitline voltage. . . . .	22
Figure6.1: Speedup with ChargeCache, NUAT and Low-Latency DRAM for single-core and eight-core workloads ((a) Single-core workloads, (b) Eight-core workloads). . . . .	27
Figure6.2: DRAM energy reduction of ChargeCache. . . . .	28
Figure6.3: ChargeCache hit rate for single-core and eight-core systems at 1ms <i>cache duration</i> . . . . .	30
Figure6.4: Speedup versus ChargeCache capacity. . . . .	31
Figure6.5: Speedup and ChargeCache hit rate for different <i>cache durations</i> . . . . .	31





## LIST OF TABLES

	<u>Page</u>
Table5.1: Simulated system configuration . . . . .	23
Table6.1: tRCD and tRAS for different <i>cache durations</i> (determined via <i>SPICE simulations</i> ) . . . . .	30





## ABBREVIATIONS

<b>DRAM</b>	: Dynamic Random-Access Memory
<b>SRAM</b>	: Static Random-access Memory
<b>DDR</b>	: Double Data-rate
<b>LLC</b>	: Last-level Cache
<b>RTL</b>	: Row-level Temporal Locality
<b>CMOS</b>	: Complementary Metal Oxide Semiconductor
<b>PCB</b>	: Printed Circuit Board
<b>I/O</b>	: Input/Output
<b>HCRAC</b>	: Highly-charged Row Address Cache



## 1. INTRODUCTION

In the last few decades, new microarchitectural techniques successfully delivered significant performance improvement to the microprocessors. At the same time, advances in the manufacturing technology, which shrunk the transistor size, provided additional processing power mainly by enabling more transistors to fit to the same die area. On the other hand, capacity of the memories also increased dramatically but the improvement in the speed of the memory was not high enough to catch up with the processors. The disparity between the performance of the processors and memory devices introduced a system-level bottleneck problem which is typically known as the "memory wall" [94, 95]. In today's multi-core era, that bottleneck is even exaggerated by the increased bandwidth requirements due to the simultaneously operating processor cores where each of them generate a significant amount of memory accesses.

DRAM technology is commonly used as the main memory of modern computer systems. This is because DRAM is at a more favorable point in the trade-off spectrum of density (cost-per-bit) and access latency compared to other technologies like SRAM or flash. However, commodity DRAM devices are heavily optimized to maximize cost-per-bit. In fact, the latency of commodity DRAM has not reduced significantly in the past decade [48, 65].

To mitigate the negative effects of long DRAM access latency, existing systems rely on several major approaches. First, they employ large on-chip caches to exploit the temporal and spatial locality of memory accesses. However, cache capacity is limited by chip area. Even caches as large as tens of megabytes may not be effective for some applications due to very large working sets and memory access characteristics that are not amenable to caching [35, 53, 69, 72, 74]. Second, systems employ aggressive prefetching techniques to preload data from memory before it is needed [2, 10, 86]. However, prefetching is inefficient for many irregular access patterns and it increases the bandwidth requirements and interference in the memory system [18, 20, 21, 43]. Third, systems employ multithreading [83, 91]. However, this approach increases contention in the memory system [14, 19, 58, 63] and does not aid single-thread performance [36, 90]. Fourth, systems exploit memory level parallelism [13, 25, 61, 63, 64]. The DRAM architecture provides various levels of parallelism that can be exploited to simultaneously process multiple memory requests generated by modern processor architectures [45, 64, 70, 92]. While prior works [15, 33, 45, 63, 68] proposed techniques to better utilize the available parallelism, the benefits of these techniques are limited due to 1) address dependencies among instructions in the programs [3, 22, 60], and 2) resource conflicts in the memory subsystem [41, 75]. Unfortunately, *none* of these four approaches *fundamentally* reduce memory latency at its source and the DRAM latency continues to be a performance bottleneck in many systems.

The latency of DRAM is heavily dependent on the design of the DRAM chip architecture, specifically the length of a wire called *bitline*. A DRAM chip consists of millions of DRAM cells. Each cell is composed of a transistor-capacitor pair. To access data from a cell, DRAM uses a component called *sense amplifier*. Each cell is connected to a sense amplifier using a *bitline*. To amortize the large cost of the sense amplifier, hundreds of DRAM cells are connected to the same bitline [48]. Longer bitlines lead to increase in resistance and parasitic capacitance on the path between the DRAM cell and the sense amplifier. As a result, longer bitlines result in higher DRAM access latency [47, 48, 85].

One simple approach to reduce DRAM latency is to use shorter bitlines. In fact, some specialized DRAM chips [26, 77, 101] offer lower latency by using shorter bitlines compared to commodity DRAM chips. Unfortunately, such chips come at a significantly higher cost as they reduce the overall density of the device because they require more sense amplifiers, which occupy significant area [48]. Therefore, such specialized chips are usually not desirable for systems that require high memory capacity [11]. Prior works have proposed several heterogeneous DRAM architectures (e.g., segmented bitlines [48], asymmetric bank organizations [85]) that divide DRAM into two regions: one with low latency, and another with slightly higher latency. Such schemes propose to map frequently accessed data to the low-latency region, thereby achieving lower average memory access latency. However, such schemes require 1) non-negligible changes to the cost-sensitive DRAM design, and 2) mechanisms to identify, map, and migrate frequently-accessed data to low-latency regions. As a result, even though they reduce the latency for some portions of the DRAM chip, they may be difficult to adopt.

**Our goal** in this work is to design a mechanism to reduce the average DRAM access latency without modifying the existing DRAM chips. We achieve this goal by exploiting two major observations we make in this thesis.

**Observation 1.** We find that, due to DRAM bank conflicts [41, 75], many applications tend to access rows that were recently closed (i.e., closed within a very short time interval). We refer to this form of temporal locality where certain rows are closed and opened again frequently as *Row Level Temporal Locality (RLTL)*. An important outcome of this observation is that a DRAM row remains in a *highly-charged* state when accessed for the second time within a short interval after the prior access. This is because accessing the DRAM row inherently replenishes the charge within the DRAM cells (just like a refresh operation does) [9, 24, 50, 51, 66, 82].

**Observation 2.** The amount of charge in DRAM cells determines the required latency for a DRAM access. If the amount of charge in the cell is low, the sense amplifier completes its operation in longer time. Therefore, DRAM access latency increases. A DRAM cell loses its charge over time and the charge is replenished by a refresh operation or an access to the row. The access latency of a cell whose charge has been replenished recently can thus be significantly lower than the access latency of a cell that has less charge.

We propose a new mechanism, called *ChargeCache* [29], that reduces average DRAM access latency by exploiting these two observations. The **key idea** is to track the addresses of recently-accessed (i.e., highly-charged) DRAM rows and serve accesses

to such rows with lower latency. Based on our observation that workloads typically exhibit significant Row-Level Temporal Locality (see Section 3), our experimental results on multi-programmed applications show that, on average, ChargeCache can reduce the latency of 67% of all DRAM row activations.

The operation of ChargeCache is straightforward. The memory controller maintains a small table that contains the addresses of a set of recently-accessed DRAM rows. When a row is evicted from the row-buffer, the address of that row, which contains highly-charged cells due to its recent access, is inserted into the table.

Before accessing a new row, the memory controller checks the table to determine if the row address is present in the table. If so, the row is accessed with low latency. Otherwise, the row is accessed with normal latency. As cells leak charge over time, ChargeCache requires a mechanism to periodically invalidate entries from the table such that only highly-charged rows remain in it. Section 4 describes the implementation of ChargeCache in detail.

Our evaluations show that ChargeCache significantly improves performance over commodity DRAM for a variety of workloads. For 8-core workloads, ChargeCache improves average workload performance by 8.6% with a hardware cost of only 5.4KB and by 10.6% with a hardware cost of 43KB. As ChargeCache can only *reduce* the latency of certain accesses, it does *not* degrade performance compared to commodity DRAM. Moreover, ChargeCache can be combined with other DRAM architectures that offer low latency (e.g., [9, 12, 41, 47, 48, 67, 78, 79, 85]) to provide even higher performance. Our estimates show that the hardware area overhead of ChargeCache is only 0.24% of a 4MB cache. Our mechanism requires no changes to DRAM chips or the DRAM interface. Section 6 describes our experimental results.

We make the following **contributions**.

- We observe that, due to bank conflicts, many applications exhibit a form of locality where recently-closed DRAM rows are accessed frequently. We refer to this as *Row Level Temporal Locality (RLTL)*(see Section 3).
- We propose an efficient mechanism, ChargeCache [29], which exploits *RLTL* to reduce the average DRAM access latency by requiring changes *only* to the memory controller. ChargeCache maintains a table of recently-accessed row addresses and lowers the latency of the subsequent accesses that hit in this table within a short time interval (see Section 4).
- We comprehensively evaluate the performance, energy efficiency, and area overhead of ChargeCache. Our experiments show that ChargeCache significantly improves performance and energy efficiency across a wide variety of systems and workloads with negligible hardware overhead (see Section 6).





## 2. BACKGROUND ON MAIN MEMORY

Memories are fundamental components used in various parts of the computer systems (e.g., register, cache, buffers, main memory, etc.). A memory system consists of multiple layers of memory units where each of these units is optimized to achieve a specific goal to converge to the ideal memory which utopically has unlimited bandwidth, zero access latency, infinite capacity, and no cost. Figure 2.1 illustrates a typical memory system that is implemented in modern computer systems. Each memory unit is scaled to indicate its actual capacity and access latency. In general, low capacity memory has lower latency compared to a memory unit with higher capacity. For example, a very limited memory resource, the register file, can typically be accessed within a single cycle. Whereas, accessing shared caches may take up to few tens of cycles to complete.

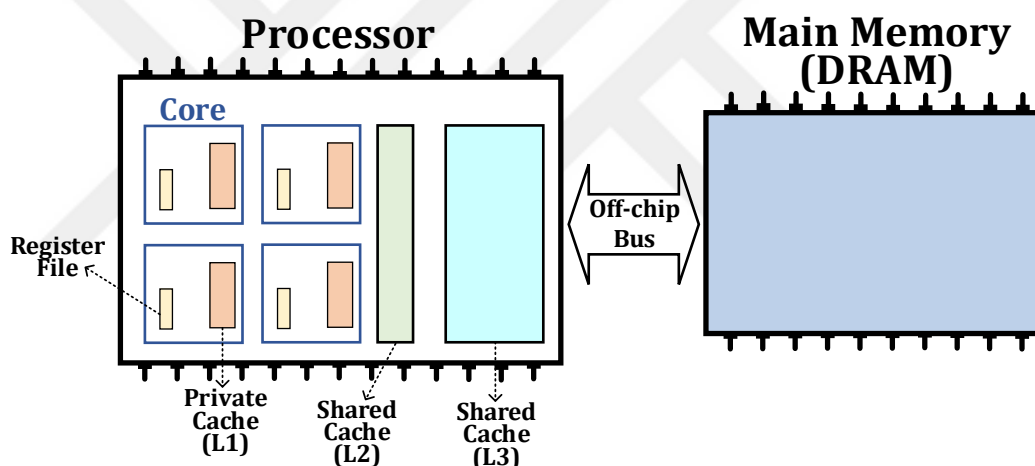


Figure 2.1: Memory system of a modern computer.

In this thesis, we mainly focus on the main memory which incurs the highest access latency in the memory system. DRAM (Dynamic Random Access Memory) technology is predominantly used as a main memory of modern system. That is because DRAM is at the most favorable point in the capacity-latency trade-off spectrum among the memory technologies that are available today. DRAM requires a special manufacturing process to benefit from its entire potential. Adapting DRAM to the common CMOS manufacturing technology, which is used to produce the processor chip (i.e, eDRAM [55]), results in higher area-per-bit usage and higher access latency compared to a custom-process DRAM. Thus, in modern systems, DRAM-based main memories are typically available as separate chip which communicates with the processor via off-chip links. Such a link imposes additional DRAM access latency.

In this section, we provide the necessary basics on DRAM organization and operation.

## 2.1 DRAM Organization

DRAM-based main memories are composed of units arranged in hierarchy of several levels (Figure 2.2). Next, we explain each level of the hierarchy in detail.

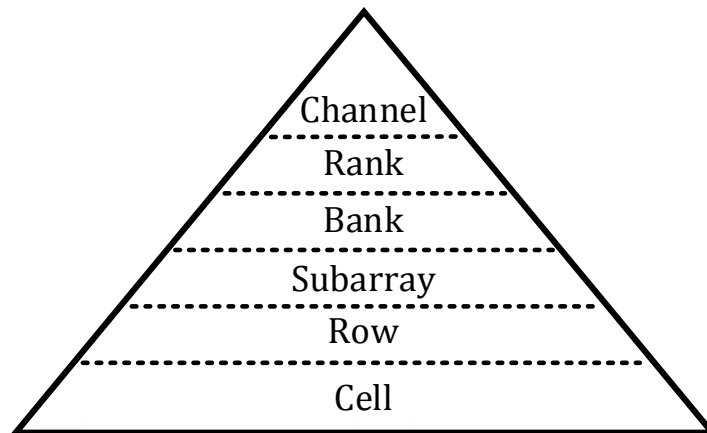


Figure 2.2: Layers of the DRAM hierarchy.

### 2.1.1 Channel

A DRAM channel is the top-level layer of the main memory hierarchy. Each channel has its own command, address, and data buses. The memory controller, a logic unit which resides inside the processor chip in modern architectures, handles the communication with the channel by issuing a set of DRAM commands to access data in the desired location (i.e., address). Figure 2.3 shows a system configuration with two memory controllers which manage a single DRAM channel each. In that particular system, the workloads running on the processor generate memory requests. A requests goes to one of the memory controller depending on the address that it targets. The address space of the system is typically spread between the two channels. Once a memory controller receives a request, it issues necessary DRAM commands to the channel to perform the access.

Several DRAM chips are put together to form a DRAM channel. In general-purpose systems (e.g., desktop computers, laptops, workstations) the chips that create a channel are soldered into a separate PCB (Printed Circuit Board) apart from the motherboard. These PCBs are called memory modules. A memory module can be directly plugged to the motherboard through the memory slots. A single channel may support one or more modules (as in Figure 2.3). If more than one modules are connected to a single channel, each module operates as a DRAM Rank which we explain next. Said that, a channel may contain one or more ranks (typically up to 4 ranks). On the other hand, in embedded systems (e.g., smartphones, single-board computers), DRAM chips are generally soldered to the motherboard along with other chips of the system.

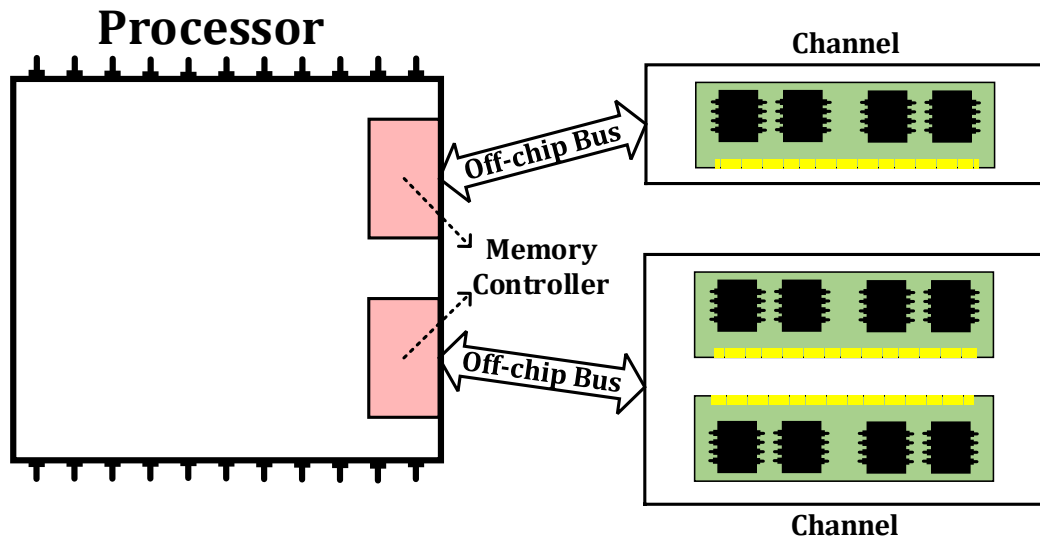


Figure 2.3: View of a system with two channels.

### 2.1.2 Rank

Different from channels, ranks do not operate in complete isolation from each other. Ranks that constitute the same channel share the address, data, and command buses (Figure 2.4). Therefore, the ranks operate in lock-step (i.e., the ranks of the same channel are time multiplexed) and do not offer pure memory access parallelism as the channels do. However, the ranks offer parallelism in lower levels of the DRAM hierarchy.

Ranks are composed of multiple DRAM chips. The number of chips depend on the data I/O width of the used chips and the width of the memory controller bus. In typically systems, the memory controller data bus is 64-bits wide. To reach the data bus width, multiple chips operate concurrently in a rank. For example, 4 DRAM chips with 16 data I/O pins each are required to form a rank.

### 2.1.3 Bank

In each rank, there are typically 8 banks available which mostly operate independently of each other. As shown in Figure 2.5 banks share the same I/O interface. They utilize that interface in lock-step fashion. The memory controller, which is on the other side of the I/O bus, can read/write to/from only one bank at once. Similarly, a data access command mostly targets a single bank. Some commands (used to initiate operations such as refresh and precharge) may apply to the all banks in a rank.

Each memory cycle, only a single bank can receive a data access command. However, since the access operation takes more than one cycle, issuing access commands to different banks consecutively enables utilization of multiple banks. For example, assume that an access takes 10 cycles to complete. After issuing an access command to the first bank, in the next cycle the memory controller may issue command to serve a request whose data is in different bank. This way, the latency of two accesses can be overlapped. Overlapping the access time of multiple requests that go to

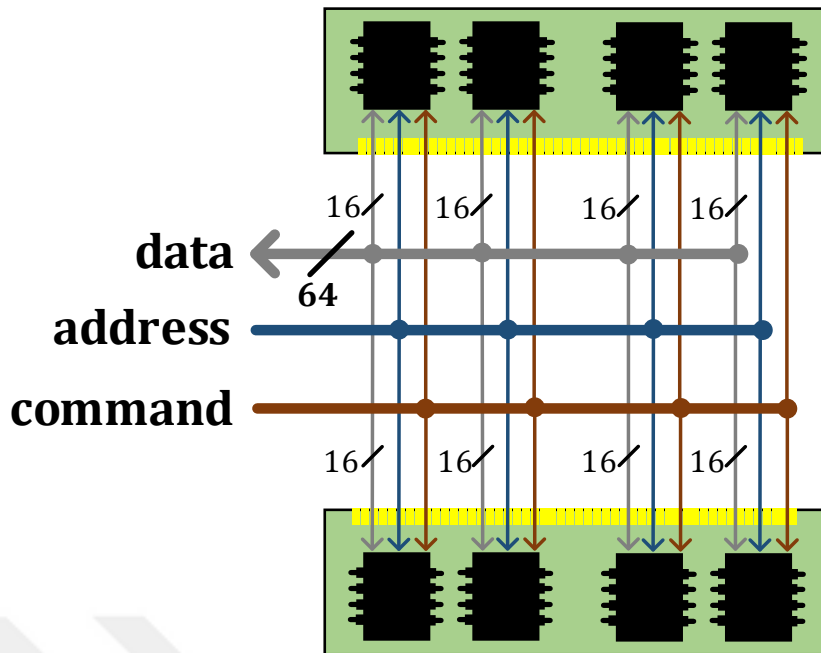


Figure 2.4: A channel which has two ranks that share data, command, and address buses.

different banks is called *Bank-Level Parallelism*. It is critical to exploit the bank-level parallelism to achieve high throughput [15, 33, 39, 40, 45, 63].

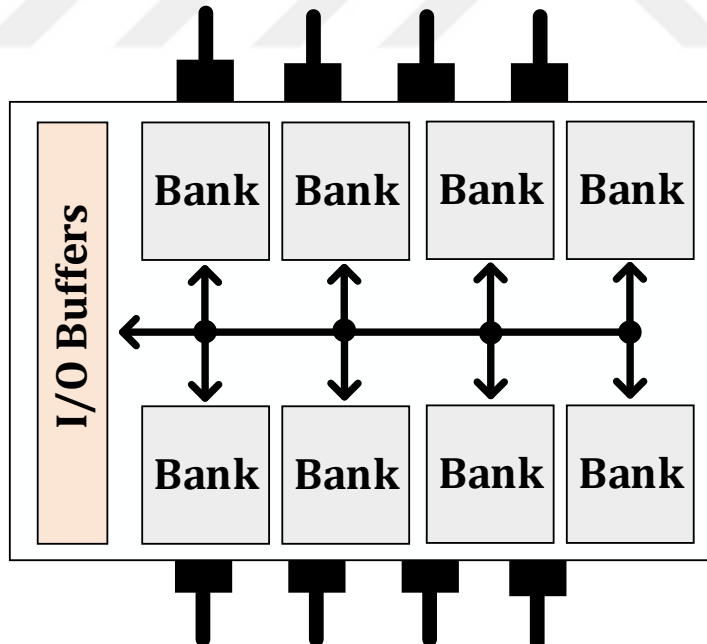


Figure 2.5: The internal structure of a rank which has 8 banks.

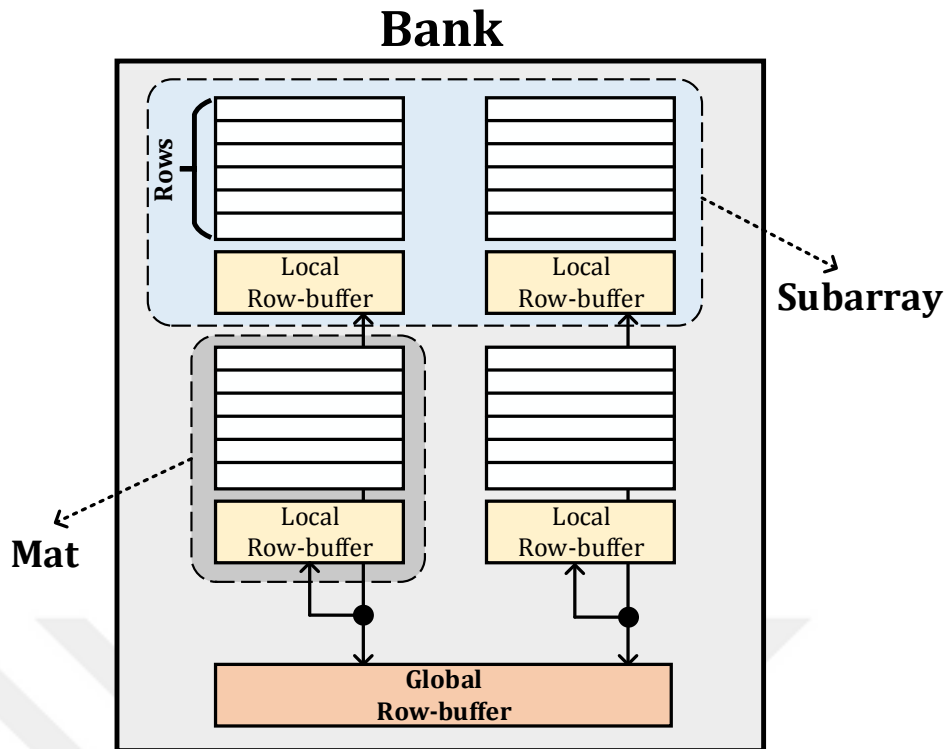


Figure 2.6: The internal structure of a bank.

#### 2.1.4 Subarray and row

Figure 2.6 depicts a DRAM bank. A bank is composed of several subarrays and a global row-buffer. Each subarray has hundreds of DRAM rows and a local row-buffer. Rows are connected to the local row-buffers via local bitlines. Similarly, local row-buffers are wired to the global row-buffer via global bitlines. The rows in a bank are grouped into subarrays to keep bitlines shorter and improve access latency by mitigating parasitic bitline capacitance. Subarrays do not provide any parallelism in current commercially available architectures. However, recent work proposes an efficient way to enable additional level of DRAM parallelism by exploiting subarray structure [41].

To perform a data access, the row that corresponds to the accessed address must be first opened by copying that row to the local row-buffer. After the data is put to the local row-buffer, the data is transferred to the global row-buffer. Opening a row is also called *Activation*. Once the data arrives the global row-buffer, the memory controller can fetch or modify a needed chunk, called column, of the global row-buffer using a single read or write command. The width of the column depends on the data I/O width of the DRAM chip.

#### 2.1.5 Cell

A DRAM cell consists of a single transistor-capacitor pair. The capacitor stores a single bit of data depending on its charge level. Asserting the wordline enables the transistor (i.e., access transistor) which couples up the capacitor and bitline. Such an operation is necessary to access a DRAM cell.

Due to the One-Transistor One-Capacitor (1T1C) architecture, a DRAM cell faces a critical leakage problem. Both the transistor and capacitor continuously leak significant amount of current which causes the DRAM cell to lose its data in milisecond-long time. As a workaround, the memory controller periodically initiates a refresh operation which restores the charge level of the cells.

## 2.2 DRAM Standards

Joint Electron Device Engineering Council (JEDEC) [100] defines standards for manufacturing a wide-range of electronic devices. JEDEC standards also involve DRAM-based memories. For example, Double Data Rate (DDR) [56] and its derivatives (such as DDR2, DDR3, DDR4) are the most widely adopted standards in DRAM memory devices. Other standards such as High Bandwidth Memory (HBM) [32], Wide I/O DRAM [17], Low-power DDR (LPDDR) [34], and Reduced Latency DRAM (RLDRAM) [101] are also available. As an example for a DRAM standard, we briefly explain the DDR3 specification which we use to evaluate our mechanism.

### 2.2.1 Double data rate type 3 (DDR3)

DDR3 standard defines a pin-interface which supports a set of commands that the memory controller uses to access (e.g., *ACT*, *PRE*, *READ*, *WRITE*) and manage (e.g., *REF*) the memory in a way we explain in Section 2.3.

DDR commands are transmitted to the DRAM module across the memory command bus. Each command is encoded using five output signals (*CKE*, *CS*, *RAS*, *CAS*, and *WE*). Enabling/disabling these signals corresponds to specific commands (as defined by the DDR standard). First, the *CKE* signal (*clock enable*) determines whether the DRAM is in “standby mode” (ready to be accessed) or “power-down mode”. Second, the *CS* (*chip selection*) signal specifies the chip that should receive the issued command. Third, the *RAS* (*row address strobe*)/*CAS* (*column address strobe*) signal is used to generating commands related to DRAM row/column operations. Fourth, the *WE* signal (*write enable*) in combination with *RAS* and *CAS*, generates the specific row/column command. For example, enabling *CAS* and *WE* together generates a *WRITE* command, while only enabling *CAS* indicates a *READ* command.

## 2.3 DDR3 Operation

DDR3 provides a set of commands which are used to perform a read/write access or other operations such as refresh. The memory controller issues these commands in specific order with certain amount of delay in between to complete the intended operation. The timing delay that must be respected between certain command is referred to as *DRAM Timing Parameters*. We explain the commands and timing parameters used to perform a typical read/write operation.

Figure 2.7 shows the different sub-steps involved in transferring the data from a DRAM cell to the sense amplifier and their mapping to DRAM commands. Each sub-step takes some time, thereby imposing some constraints (i.e., timing parameters) on when the

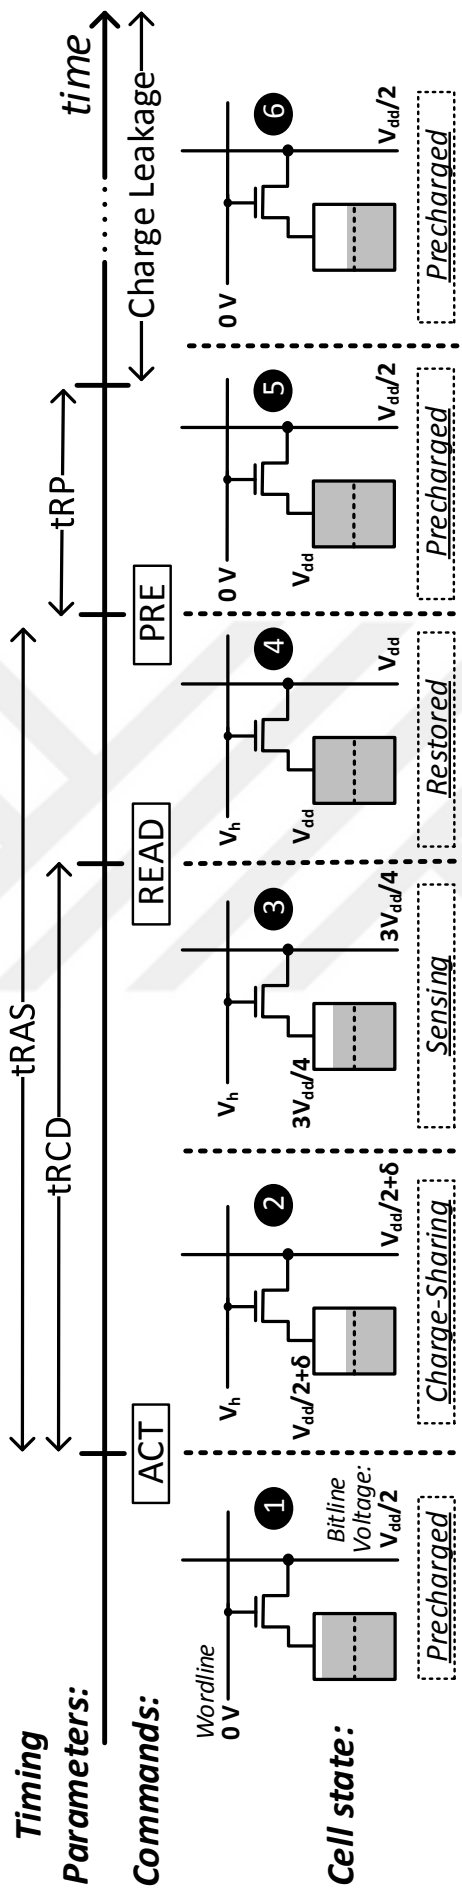


Figure 2.7: Commands that are used to read data from DRAM and the timing parameters associated with them

memory controller can issue different commands. The figure also shows the major timing parameters that govern regular DRAM operation.

In the initial *precharged* state ❶, the bitline is precharged to a voltage level of  $V_{dd}/2$ . The wordline is lowered (i.e., at 0V) and hence, the bitline is *not* connected to the capacitor. An access to the cell is triggered by the *ACT* command to the corresponding row. This command first raises the wordline (to voltage level  $V_h$ ), thereby connecting the capacitor to the bitline. Since the capacitor (in this example) is at a higher voltage level than the bitline, charge flows from the capacitor to the bitline, thereby raising the voltage level on the bitline to  $V_{dd}/2 + \delta$  ❷. This phase is called *charge sharing*. After the charge sharing phase, the sense amplifier is enabled and it detects the deviation on the bitline, and amplifies the deviation. This process, known as *sense amplification*, drives the bitline and the cell to the voltage level corresponding to the original state of the cell ( $V_{dd}$  in this example). Once the sense amplification has sufficiently progressed ❸, the memory controller can issue a *READ* or *WRITE* command to access the data from the cell. The time taken by the cell to reach this state ❸ after the *ACT* command is specified by the timing constraint *tRCD*. Once the sense amplification process is complete ❹, the bitline and the cell are both at a voltage level of  $V_{dd}$ . In other words, the original charge level of the cell is *fully-restored*. The time taken for the cell to reach this state ❹ after the *ACT* is specified by the timing constraint *tRAS*. In this state, the bitline can be precharged using the *PRE* command to prepare it for accessing a different row. This process first lowers the wordline, thereby disconnecting the cell from the bitline. It then precharges the bitline to a voltage level of  $V_{dd}/2$  ❺. The time taken for the precharge operation is specified by the timing constraint *tRP*.

**DRAM Charge Leakage and Refresh.** As DRAM cells are not ideal, they leak charge after the precharge operation [50, 51]. This is represented in state ❻ of Figure 2.7. As described in the previous section, an access to a DRAM cell fully restores the charge on the cell (see states ❹ and ❺). However, if a cell is not accessed for a sufficiently long time, it may lose too much charge that its last cell state may be flipped. To avoid such cases, DRAM cells are periodically refreshed by the memory controller using the refresh (*REF*) command. The interval at which DRAM cells should be refreshed by the controller is referred to as the *refresh interval*.

## 2.4 Memory Controller

The Memory Controller sits between the Last-level Cache (LLC) and the DRAM. Today, the memory controller is typically employed in the same chip with the processor logic, as shown in Figure 2.8.

The memory controller is mainly responsible for handling the load/store requests generated by the LLC. The bottom part of Figure 2.8 shows an illustration of functional building block of a memory controller. Due to cache misses or dirty data evictions, LLC generates load/store requests. Once received, the memory controller stores these requests in the *Request Buffer*. Then the *scheduling logic* decides which request from the request buffer to serve first. The scheduling logic (or simply scheduler) makes this decision based on a set of heuristics which may improve average request serving time (latency), fairness, or throughput. Once the scheduler makes its decision, based on the state of the target bank, the *Command Generator*



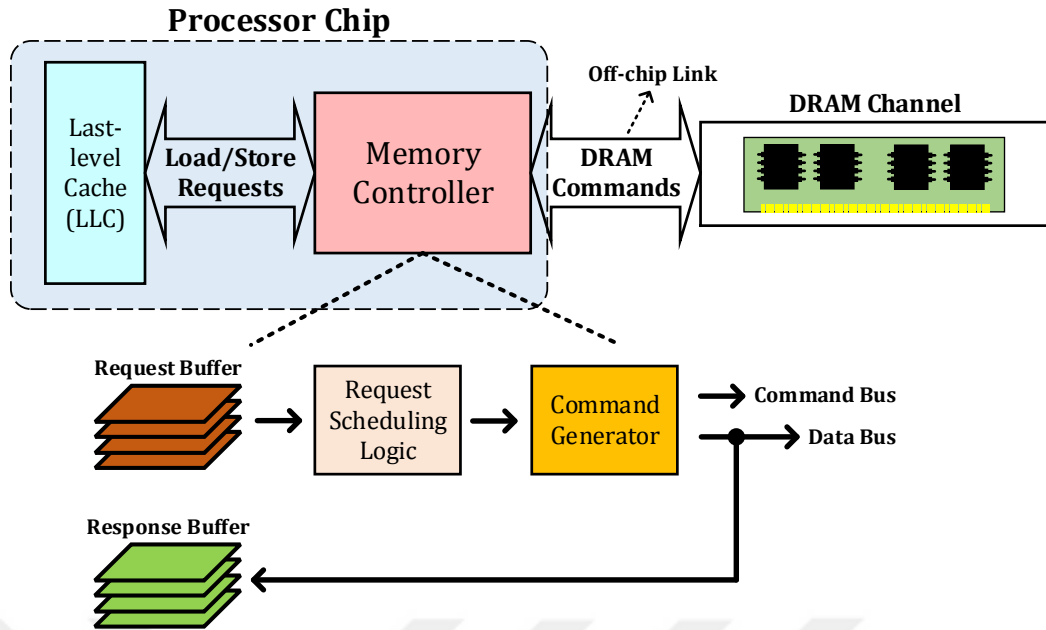


Figure 2.8: Overview of a typical memory controller.

cracks the request into appropriate DRAM commands. For instance, if the target bank has an open row and the address of that row is the same as the target row of the requests, then the command generator only issues a *READ* or *WRITE* command to the target bank. Whereas, if we have row conflict (i.e., if the address of the target row is different from the open row address) the memory controller first issues a *PRE* command to close the conflicting row. Then, by issuing an *ACT*, the memory controller activates the target row of the request that is being serviced. Thus, the output of the command generator not only depends on the decision of the scheduler, but also on the internal state of the DRAM. The memory controller also receives data from the DRAM and forwards it to the LLC to respond to the load request.

A memory controller employs smart scheduling algorithms to (i) reduce access latency, (ii) improve throuput, or (iii) provide better quality of service (QoS) among concurrently running workloads. A large number of prior work studies scheduling algorithms to improve these three aspects [1, 9, 11, 16, 19, 31, 33, 37, 39–41, 63, 66, 76, 97].



### 3. MOTIVATION

The key takeaway from DRAM operation that we exploit in this work is the fact that *cells closer to the fully-charged state can be accessed with lower activation latency* (i.e., lower  $t_{RCD}$  and  $t_{RAS}$ ) than standard DRAM specification. A recent work [82] exploits this observation to access rows, that were recently recharged via a *refresh* operation, with lower latency. Specifically, when a row needs to be activated, the memory controller determines when the row was last refreshed. If the row was refreshed recently (e.g., within  $8ms$ ), the controller uses a lower  $t_{RCD}$  and  $t_{RAS}$  for the activation.

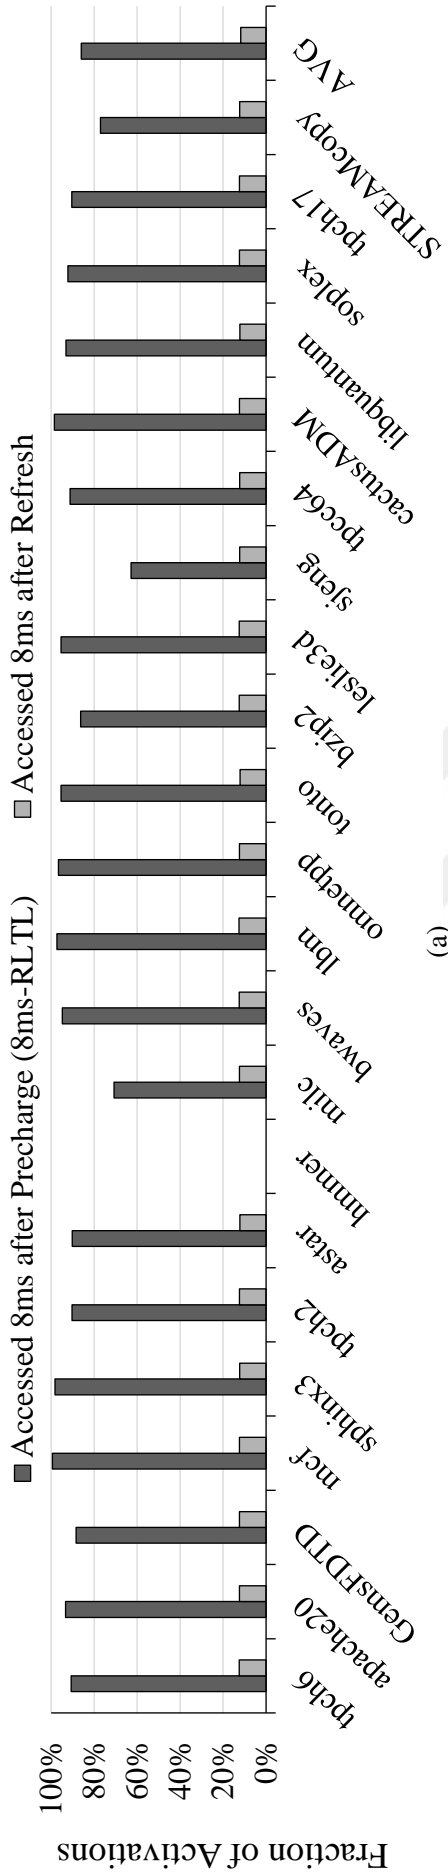
However, this refresh-based approach for lowering latency has two shortcomings. First, with the standard refresh mechanism, the refresh schedule used by the memory controller has no correlation with the memory access characteristics of the application. Therefore, depending on the point when the program begins execution, a particular row activation, due to a memory access initiated by the program, may or may not be to a recently-refreshed row. Therefore, a mechanism that reduces latency to recently-refreshed rows cannot provide consistent performance improvement. Second, if we use only the time from the last refresh to identify rows that can be accessed with low latency (i.e., highly-charged rows), we find that only 12% of all memory accesses benefit from low latency (see Figure 3.1). However, as we show next, a much greater number of rows can actually be accessed with low latency.

As we described in Section 2.3, an access to a row fully recovers the charge of its cells. Therefore, if a row is *activated twice in a short interval*, the second activate can be served with lower latency as the cells of that row would still be highly charged. We refer to this notion of row activation *locality* as *Row-Level Temporal Locality* (RLTL). We define  $t$ -RLTL of an application for a given time interval  $t$  as the fraction of row activations in which the activation occurs within the time interval  $t$  after a previous precharge to the same row. (Recall that, a row starts leaking charge only after the precharge operation as shown in Section 2.3).

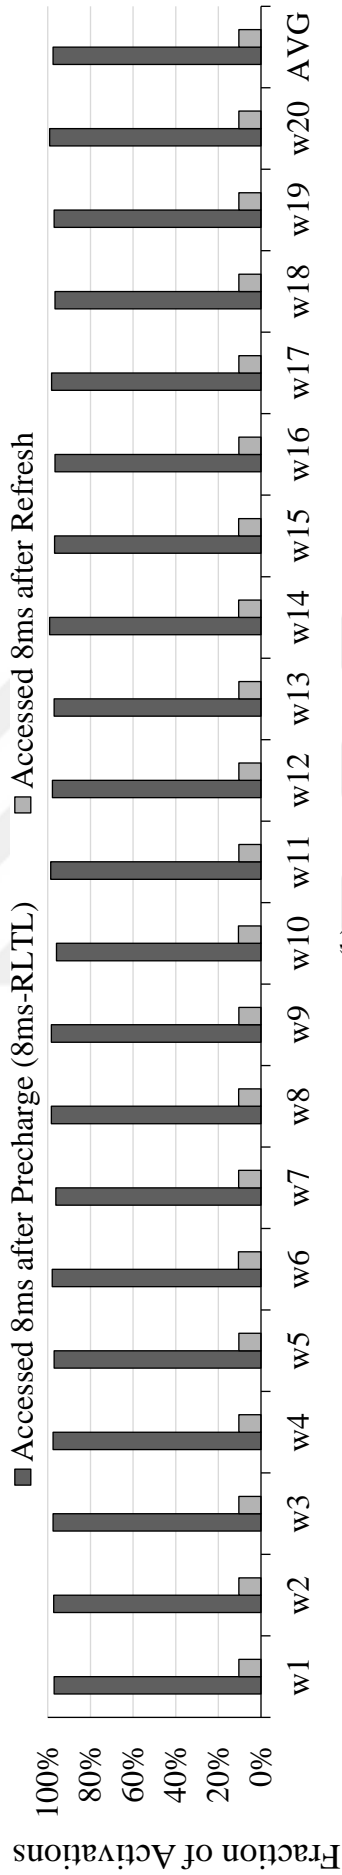
To this end, we would like to understand what fraction of rows exhibit RLTL, and thus can be accessed with low latency after a precharge operation to the row due to program behavior versus what fraction of rows are accessed soon after a refresh to the row and thus can be accessed with low latency due to a recent preceding refresh. Figure 3.1a compares the fraction of row activations of that happen within  $8ms$  after the corresponding row is refreshed to the  $8ms$ -RLTL of various applications. As shown in the figure, with the exception of *hammer*<sup>1</sup>, the  $8ms$ -RLTL (86% on average) is significantly higher than the fraction of row activations within  $8ms$  after the refresh of the row (12% on average). Figure 3.1b plots the corresponding values on an 8-core system that executes 20 multiprogrammed workloads, with randomly chosen

---

<sup>1</sup>*hammer* effectively uses the on-chip cache hierarchy. Therefore, we do not observe any requests to the main memory.



(a)



(b)

Figure 3.1: Fraction of row activations that happen 8ms after precharge (8ms-RLTL) or refresh of the row ((a) Single-core workloads, (b) Eight-core workloads).

applications for each workload. As shown, the fraction of row activations within  $8ms$  after refresh is almost the same as that of the single-core workloads. This is because the refresh schedule has no correlation with the application access pattern. On the other hand, the  $8ms$ -RLTL for the 8-core workloads is much higher than that of the single-core workloads. This is because, in multi-core systems, the exacerbated bank-level contention [40, 45, 59, 62, 63, 97] results in row conflicts, which in turn results in rows getting closed and activated within shorter time intervals, leading to a high RLTL.

Figure 3.2 shows the RLTL for different single-core and 8-core workloads with five different time intervals (from  $0.125ms$  to  $32ms$ ) as a stacked bar and two different DRAM row management policies, namely, open-row and closed-row [1, 39]. For each workload, the first bar represents the results for the open-row policy, and the second bar represents the results for the closed-row policy. The open-row policy prioritizes row-buffer hits by keeping the row open until a request to another row is scheduled (bank conflict). In contrast, the closed-row policy proactively closes the active row after servicing all row-hit requests in the request buffer.

For single-core workloads (Figure 3.2a), regardless of the row-buffer policy, even the average  $0.125ms$ -RLTL is 66%. In other words, 66% of all the row activations occur within  $0.125ms$  after the row was previously precharged. For 8-core workloads (Figure 3.2b), due to the additional bank conflicts, the average  $0.125ms$ -RLTL is 77%, significantly higher than that for the single-core workloads. Similar to the single-core workloads, the row-buffer policy does not have a significant impact on the RLTL for the 8-core workloads.

**Key Observation and Our Goal.** We observe that *many* applications exhibit high row-level temporal locality. In other words, for many applications, a significant fraction of the row activations occur within a small interval after the corresponding rows are precharged. As a result, such row activations can be served with lower activation latency than specified by the DRAM standard. **Our goal** in this work is to exploit this observation to reduce the effective DRAM access latency by tracking recently-accessed DRAM rows in the memory controller and reducing the latency for their next access(es). To this end, we propose an efficient mechanism, ChargeCache, which we describe in the next section.

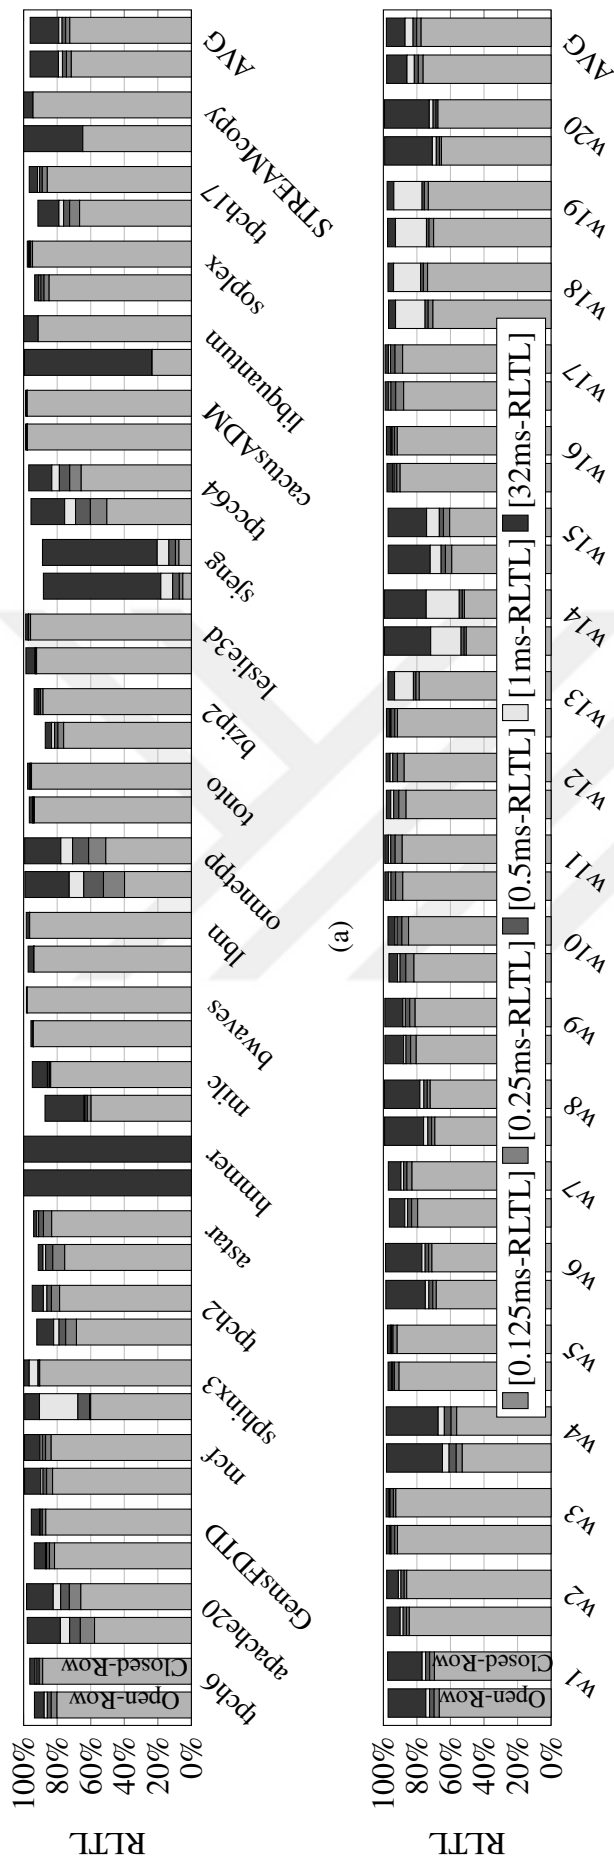


Figure 3.2: RTTL for various time intervals ((a) Single-core workloads, (b) Eight-core workloads).

## 4. CHARGECACHE

ChargeCache is based on three observations: 1) rows that are highly-charged can be accessed with lower activation latency, 2) activating a row refreshes the charge on the cells of that row and the cells start leaking only after the following precharge command, and 3) many applications exhibit high row-level temporal locality, i.e., recently-activated rows are more likely to be activated again. Based on these observations, ChargeCache tracks rows that are recently activated, and serves future activations to such rows with lower latency by lowering the DRAM timing parameters for such activations.

### 4.1 High-level Overview

At a high level, ChargeCache adds a small table (or cache) to the memory controller that tracks the addresses of recently-accessed DRAM rows, i.e., highly-charged rows. ChargeCache performs three operations. First, when a precharge command is issued to a bank, ChargeCache inserts the address of the row that was activated in the corresponding bank to the table (Section 4.2.1). Second, when an activate command is issued, ChargeCache checks if the corresponding row address is present in the table. If the address is not present, then ChargeCache uses the standard DRAM timing parameters to issue subsequent commands to the bank. However, if the address of the activated row is present in the table, ChargeCache employs reduced timing parameters for subsequent commands to that bank (Section 4.2.2). Third, ChargeCache invalidates entries from the table to ensure that rows corresponding to valid entries can indeed be accessed with lower access latency (Section 4.2.3).

We named the mechanism *ChargeCache* as it provides a *cache*-like benefit, i.e., latency reduction based on a locality property (i.e., RLTL), and does so by taking advantage of the *charge* level stored in a recently-activated row. The mechanism could potentially be used with current and emerging DRAM-based memories where the stored charge level leads to different access latencies. We explain how ChargeCache can be applied to other DRAM standards in Section 7.2.

In the following section, we describe the different components and operation of ChargeCache in more detail. In Section 4.3, we present the results of our SPICE simulation that analyzes the potential latency reduction that can be obtained using ChargeCache.

### 4.2 Detailed Design

ChargeCache adds two main components to the memory controller. Figure 4.1 highlights these components. The first component is a tag-only cache that stores the addresses of a subset of highly-charged DRAM rows. We call this cache the

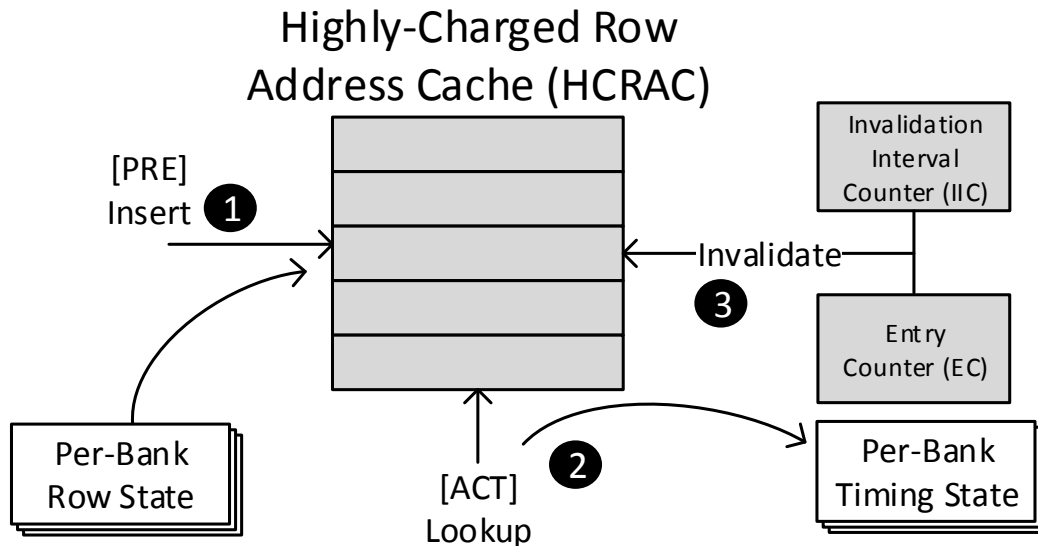


Figure 4.1: Components of the ChargeCache Mechanism

*Highly-Charged Row Address Cache (HCRAC)*. We organize HCRAC as a set-associative structure similar to the processor caches. The second component is a set of two counters that ChargeCache uses to invalidate entries from the HCRAC that can potentially point to rows that are no longer highly-charged. As described in the previous section, there are three specific operations with respect to ChargeCache: 1) insert, 2) lookup, and 3) invalidate. We now describe these operations in more detail.

#### 4.2.1 Inserting rows into HCRAC

When a *PRE* command is issued to a bank, ChargeCache inserts the address of the row that was activated in the corresponding bank into the HCRAC ❶. Although the *PRE* command itself is associated only with the bank address, the memory controller has to maintain the address of the row that is activated in each bank (if any row is activated) so that it can issue appropriate commands when a bank receives a memory request. ChargeCache obtains the necessary row address information directly from the memory controller. Some DRAM interfaces [56] allow the memory controller to precharge all banks with a single command. In such cases, ChargeCache inserts the addresses of the activated rows across *all* the banks into the HCRAC.

Just like any other cache, HCRAC contains a limited number of entries. As a result, when a new row address is inserted, ChargeCache may have to evict an already valid entry from the HCRAC. While such evictions can potentially result in wasted opportunity to reduce DRAM latency for some row activations, our evaluations show that even with a small HCRAC (e.g., 128-entries), ChargeCache can provide significant performance improvement (see Section 6).



### 4.2.2 Employing lowered DRAM timing constraints

To employ lower latency for highly-charged rows, the memory controller maintains two sets of timing constraints, one for regular DRAM rows, and another for highly-charged DRAM rows. While we evaluate the potential reduction in timing constraints that can be enabled by ChargeCache, we expect the lowered timing constraints for highly-charged rows to be part of the standard DRAM specification.

On each *ACT* command, ChargeCache looks up the corresponding row address in the HCRAC ②. Upon a hit, ChargeCache employs lower  $t_{RCD}$  and  $t_{RAS}$  for the subsequent *READ/WRITE* and *PRE* operations, respectively. Upon a miss, ChargeCache employs the default timing constraints for the subsequent commands.

### 4.2.3 Invalidating stale rows from HCRAC

Unlike conventional caches, where an entry can stay valid as long as it is not explicitly evicted, entries in HCRAC have to be invalidated after a specific time interval. This is because as DRAM cells continuously leak charge, a highly-charged row will no longer be highly-charged after a specific time interval.

One simple way to invalidate stale entries would be to use a clock to track time and associate each entry with an expiration time. Upon a hit in the HCRAC, ChargeCache can check if the entry is past the expiration time to determine which set of timing parameters to use for the corresponding row. However, this scheme increases the storage cost and complexity of implementing ChargeCache.

We propose a simpler, periodic invalidation scheme that is similar to how the memory controller issues refresh commands [51]. Our mechanism uses two counters, namely, the *Invalidation Interval Counter* (IIC) and the *Entry Counter* (EC). We assume that the HCRAC contains  $k$  entries and the number of processor cycles for which a DRAM row stays highly-charged after a precharge is  $C$ . IIC cyclically counts up to  $C/k$ , and EC cyclically counts up to  $k$ . Initially, both IIC and EC are initialized to zero. IIC is incremented every cycle. Whenever IIC reaches  $C/k$ , 1) the entry in the HCRAC pointed to by EC is invalidated, 2) EC is incremented, and 3) IIC is cleared. Whenever EC reaches  $k$ , it is cleared. This mechanism invalidates every entry in the HCRAC once every  $C$  processor cycles. Therefore, it ensures that any valid entry in the HCRAC indeed corresponds to a highly-charged row. While our mechanism can prematurely invalidate an entry, our evaluations show that the loss in performance benefit due to such premature evictions is negligible.

## 4.3 Reduction in DRAM Timing Parameters

We evaluate the potential reduction in  $t_{RCD}$  and  $t_{RAS}$  for ChargeCache using circuit-level SPICE simulations. We implement the DRAM sense amplifier circuit using 55nm DDR3 model parameters [103] and PTM low-power transistor models [98, 102]. Figure 4.2 plots the variation in bitline voltage level during cell activation for different initial charge amounts of the cell.

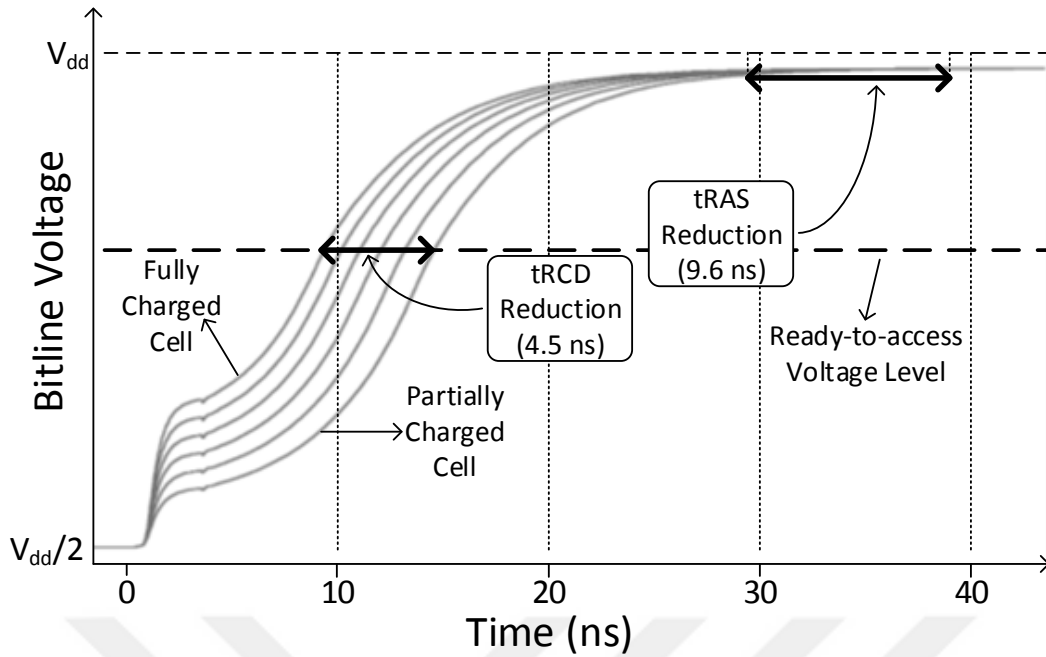


Figure 4.2: Effect of initial cell charge on bitline voltage.

Depending on the initial charge (i.e., voltage level) of the cell, the bitline voltage increases at different speeds. When the cell is *fully-charged*, the sense amplifier is able to drive the bitline voltage to the *ready-to-access voltage level* in only  $10\text{ns}$ . However, a partially-charged cell (i.e., one that has not been accessed for  $64\text{ms}$ ) brings the bitline voltage up slower. Specifically, the bitline connected to such a partially-charged cell reaches the ready-to-access voltage level in  $14.5\text{ns}$ . Since DRAM timing parameters are dictated by this worst-case partially-charged state right before the refresh interval, we can achieve  $4.5\text{ns}$  reduction in  $t_{\text{RCD}}$  for a fully-charged cell. Similarly, the charge of the cell capacitor is restored at different times depending on the initial voltage of the cell. For a fully-charged cell, this results in  $9.6\text{ns}$  reduction in  $t_{\text{RAS}}$ .

In practice, we expect the DRAM manufacturers to identify the lowered timing constraints for different caching durations. Today, DRAM manufacturers test each DRAM chip to determine if it meets the timing specifications. Similarly, we expect the manufacturers would also test each chip to determine if it meets the ChargeCache timing constraints.

*Caching duration* (i.e., how long a row address stays in ChargeCache) provides a trade-off between ChargeCache hit-rate and the DRAM access latency reduction. A longer *caching duration* leads to a longer *Invalidation Interval*. Thus, a row address stays a longer time in ChargeCache. This creates an opportunity to increase ChargeCache hit-rate. On the other hand, with a longer *caching duration*, the amount of charge that remains in DRAM cells at the end of the duration decreases. Consequently, the room for reducing  $t_{\text{RCD}}$  and  $t_{\text{RAS}}$  shrinks. As Figure 3.1 indicates a very high *RLTL* even with a  $0.125\text{ms}$  duration, we believe sacrificing ChargeCache hit-rate for DRAM access latency is a reasonable design choice. Therefore, we assume a  $1\text{ms}$  *caching duration* and a corresponding  $4/8$  cycle reduction in  $t_{\text{RCD}}/t_{\text{RAS}}$  (determined using SPICE simulations) for a DRAM bus clocked at  $800\text{MHz}$  frequency. To support our design decision, we also analyze the effect of various *caching durations* in Section 6.4.2.

## 5. METHODOLOGY

To evaluate the performance of ChargeCache, we use a cycle-accurate DRAM simulator, Ramulator [42, 104], in CPU-trace-driven mode. CPU traces are collected using a Pintool [54]. Table 5.1 lists the configuration of the evaluated systems. We implement the HCRAC similarly to a 2-way associative cache that uses the LRU policy.

Table 5.1: Simulated system configuration

Processor	1-8 cores, 4GHz clock frequency, 3-wide issue, 8 MSHRs/core, 128-entry instruction window
Last-level Cache	64B cache-line, 16-way associative, 4MB cache size
Memory Controller	64-entry read/write request queues, FR-FCFS scheduling policy [76, 99], open/closed row policy [39, 40] for single/multi core
DRAM	DDR3-1600 [56], 800MHz bus frequency, 1/2 channels, 1 rank/channel, 8 banks/rank, 64K rows/bank, 8KB row-buffer size, tRCD/tRAS 11/28 cycles
ChargeCache	128-entry (672 bytes)/core, 2-way associativity, LRU replacement policy, 1ms caching duration, tRCD/tRAS reduction 4/8 cycles

For area, power, and energy measurements, we modify McPAT [49] to implement ChargeCache using 22nm process technology. We also use DRAMPower [7] to obtain power/energy results of the off-chip main memory subsystem. We feed DRAMPower with DRAM command traces obtained from our simulations using Ramulator.

We run 22 workloads from SPEC CPU2006 [105], TPC [107] and STREAM [106] benchmark suites. We use SimPoint [28] to obtain traces from representative phases of each application. For single-core evaluations, unless stated otherwise, we run each workload for 1 billion instructions. For multi-core evaluations, we use 20 multi-programmed workloads by assigning a randomly-chosen application to each core. We evaluate each configuration with its best performing row-buffer management policy. Specifically, we use the open-row policy for single-core and closed-row policy for multi-core configurations. We simulate the benchmarks until each core executes at least 1 billion instructions. For both single and multi-core configurations, we first warm up the caches and ChargeCache by fast-forwarding 200 million cycles.

We measure performance improvement for single-core workloads using the Instructions per Cycle (IPC) metric. We measure multi-core performance using the

weighted speedup [84] metric. Prior work has shown that weighted speedup is a measure of system throughput [23].



## 6. EVALUATION

We experimentally evaluate the following mechanisms: 1) ChargeCache [29], 2) NUAT [82], which accesses *only* rows that are *recently-refreshed* at lower latency than the DRAM standard, 3) ChargeCache + NUAT, which is a combination of ChargeCache and NUAT [82] mechanisms, and 4) Low-Latency DRAM (LL-DRAM) [26], which is an idealized comparison point where we assume *all* rows in DRAM can be accessed with low latency, compared to our baseline DDR3-1600 [56] memory, at any time, regardless of when they are accessed or refreshed.

We primarily use a 128-entry ChargeCache, which provides an effective trade-off between performance and hardware overhead. We analyze sensitivity to ChargeCache capacity in Section 6.4.1. We evaluate LL-DRAM to show the upper limit of performance improvement that can be achieved by reducing  $tRCD$  and  $tRAS$ . LL-DRAM uses, for all DRAM accesses, the same reduced values for these timing parameters as we use for ChargeCache hits. In other words, LL-DRAM is the same as ChargeCache with a 100% hit rate.

We compare the performance of our mechanism against the most closely related previous work, NUAT [82], and also show the benefit of using both ChargeCache and NUAT together. The key idea of NUAT is to access *recently-refreshed* rows at low latency, because these rows are already highly-charged. Thus, NUAT does not usually access rows that are *recently-accessed* at low latency, and hence it does not exploit existing RLTL (Row-Level Temporal Locality) present in many applications. As we show in Section 3, the fraction of activations that are to rows that are recently-accessed by the application is much higher than the fraction of activations that are to rows that are recently-refreshed. In other words, many workloads have very high RLTL, which is not exploited by NUAT. As a result, we expect ChargeCache to significantly outperform NUAT since it can reduce DRAM latency for a much greater fraction of DRAM accesses than NUAT. To quantitatively prove our expectation that ChargeCache should widely outperform NUAT, we implement NUAT in Ramulator using the default 5PB configuration used in [82].

Note that NUAT bins the rows into different latency categories based on how recently they were refreshed. For instance, NUAT accesses rows that were refreshed between 0 – 6ms ago with different  $tRCD$  and  $tRAS$  parameters than rows that were refreshed between 6 – 16ms ago. We determined the different timing parameters of different NUAT bins using SPICE simulations. Although ChargeCache can implement a similar approach to NUAT by using multiple *caching durations*, our *RLTL* results motivate a single *caching duration* since a row is typically accessed within 1ms (as shown in Section 3). A row that hits in ChargeCache is always accessed with reduced timings (Section 4.3).

## 6.1 Impact on Performance

Figure 6.1 shows the performance of single-core and eight-core workloads. The figure also includes the number of row misses per kilo-cycles (RMPKC) to show row activation intensity, which provides insight into the RLTL of the workload.

**Single-core.** Figure 6.1a shows the performance improvement over the baseline system for single-core workloads. These workloads are sorted in ascending order of RMPKC. ChargeCache achieves up to 9.3% (an average of 2.1%) speedup.

Our mechanism outperforms NUAT and achieves a speedup close to LL-DRAM with a few exceptions. Applications that have a wide gap in performance between ChargeCache and LL-DRAM (such as *mcf*, *omnetpp*) access a large number of DRAM rows and exhibit high row-reuse distance [37]. A high row-reuse distance indicates that there is large number of accesses to other rows between two accesses to the same row. Due to this reason, ChargeCache *cannot* retain the addresses of highly-charged rows until the next access to that row. Increasing the number of ChargeCache entries or employing cache management policies aware of reuse distance or thrashing [16, 72, 81] may improve the performance of ChargeCache for such applications. We leave the evaluation of these methods for future work. We conclude that ChargeCache significantly reduces execution time for most high-RMPKC workloads and outperforms NUAT for all but few workloads.

**Eight-core.** Figure 6.1b shows the speedup on eight-core multiprogrammed workloads. On average, ChargeCache and NUAT improve performance by 8.6% and 2.5%, respectively. Employing ChargeCache in combination with NUAT achieves a 9.6% speedup, which is only 3.8% less than the improvement obtained using LL-DRAM. Although the multiprogrammed workloads are composed of the *same* applications as in single-core evaluations, we observe much higher performance improvements among eight-core workloads. The reason is twofold.

First, since multiple cores share a limited capacity LLC, simultaneously running applications compete for the LLC. Thus, individual applications access main memory more often, which leads to higher RMPKC. This makes the workload performance more sensitive to main memory latency [5, 31, 41]. Second, the memory controllers receive memory requests from multiple simultaneously-running applications to a limited number of memory banks. Such requests are likely to target different rows since they use separate memory regions and these regions map to separate rows. Therefore, applications running concurrently exacerbate the bank-conflict rate and increase the number of row activations that hit in ChargeCache.

Overall, ChargeCache improves performance by up to 8.1% (11.3%) and 2.1% (8.6%) on average for single-core (eight-core) workloads. It outperforms NUAT for most of the applications and using NUAT in combination with ChargeCache improves the performance slightly further.

## 6.2 Impact on DRAM Energy

ChargeCache incurs negligible area and power overheads (Section 6.3). Because it reduces execution time with negligible overhead, it leads to significant energy

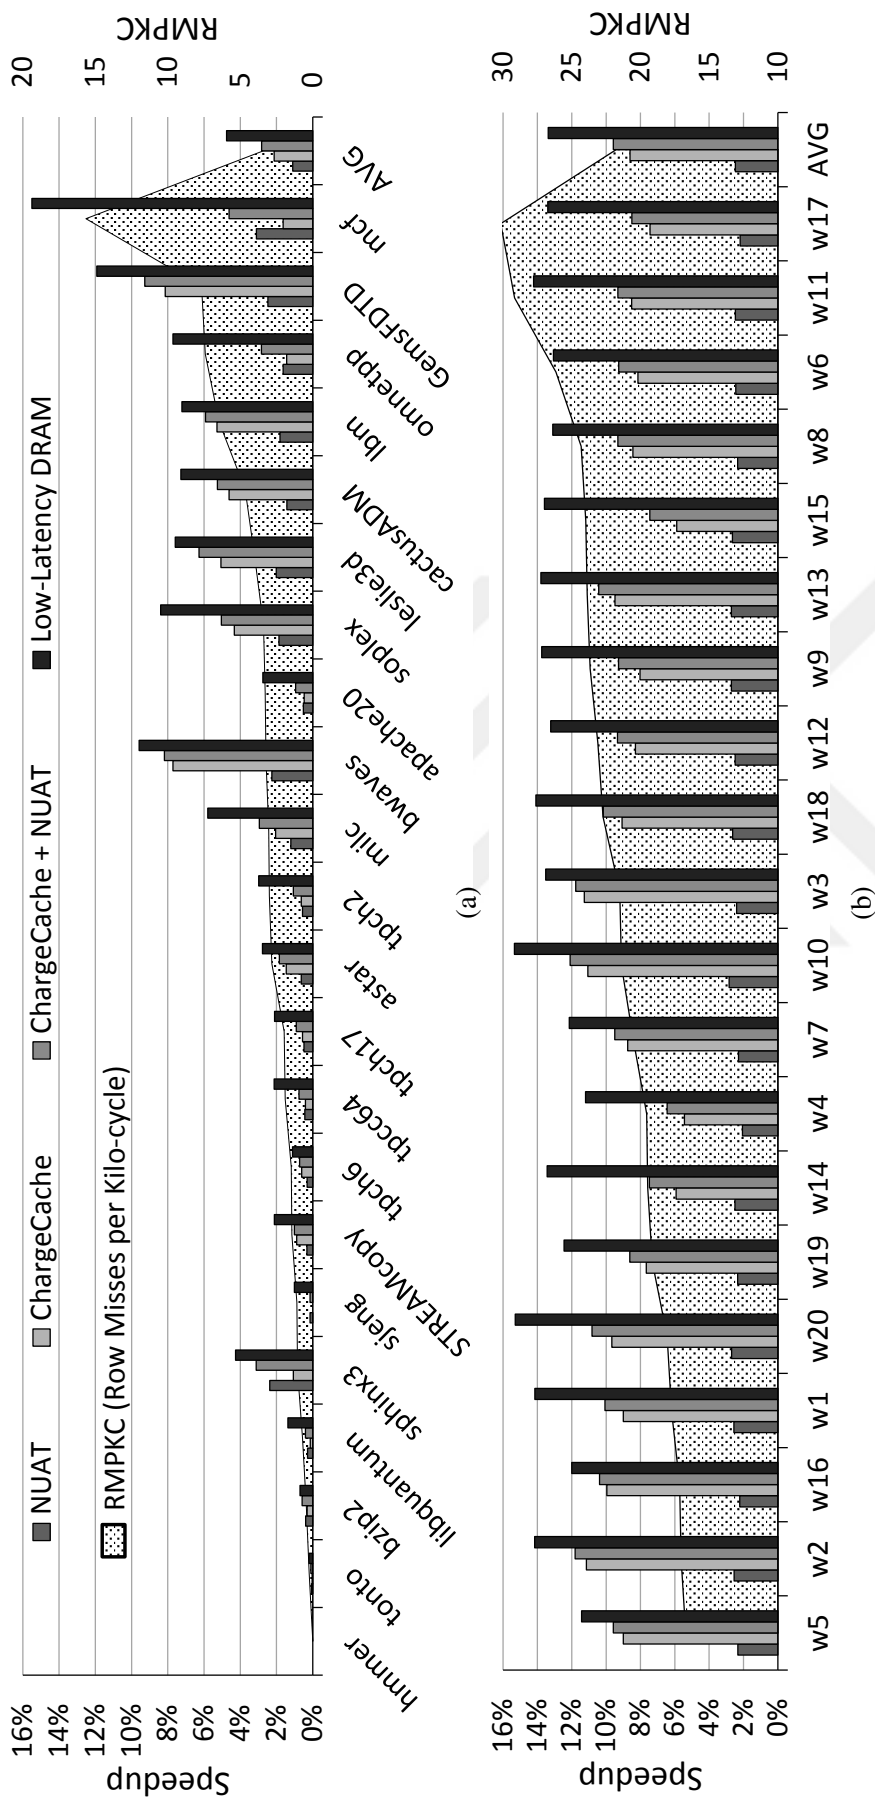


Figure 6.1: Speedup with ChargeCache, NUAT and Low-Latency DRAM for single-core and eight-core workloads ((a) Single-core workloads, (b) Eight-core workloads).

savings. Even though ChargeCache increases the energy efficiency of the entire system, we quantitatively evaluate the energy savings only for the DRAM subsystem since Ramulator [42] does not have a detailed CPU model.

Figure 6.2 shows the average and maximum DRAM energy savings for single-core and eight-core workloads. ChargeCache reduces energy consumption by up to 6.9% (14.1%) and on average 1.8% (7.9%) for single-core (eight-core) workloads. We conclude that ChargeCache is effective at improving the energy efficiency of the DRAM subsystem, as well as the entire system.

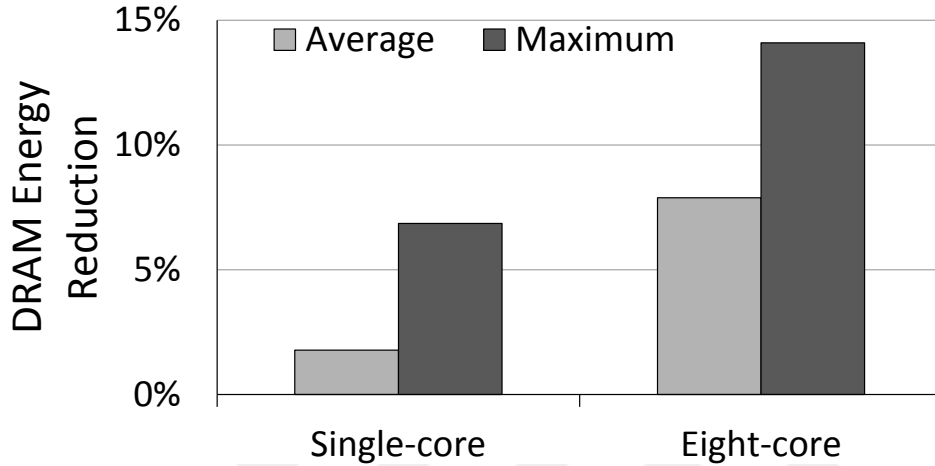


Figure 6.2: DRAM energy reduction of ChargeCache.

### 6.3 Area and Power Consumption Overhead

HCRAC (Highly-Charged Row Address Cache) is the most area/power demanding component of ChargeCache. The overhead of *EC* and *IIC* is negligible since they are just two simple counters. As we replicate ChargeCache on a per-core and per-memory channel basis, the total area and power overhead ChargeCache introduces depends on the number of cores and memory channels.<sup>2</sup> The total storage requirement is given by Equation 6.1, where *C* and *MC* are the number of cores and memory channels, respectively. *LRUbits* depends on ChargeCache associativity. *EntrySize* is calculated using Equation 6.2, where *R*, *B*, and *Ro* are the number of ranks, banks, and rows in DRAM, respectively.

$$Storage_{bits} = C * MC * Entries * (EntrySize_{bits} + LRU_{bits}) \quad (6.1)$$

$$EntrySize_{bits} = \log_2(R) + \log_2(B) + \log_2(Ro) + 1 \quad (6.2)$$

**Area.** Our eight-core configuration has two memory channels. This introduces a total of 5376 bytes in storage requirement for a 128-entry ChargeCache, corresponding to an area of 0.022 mm<sup>2</sup>. This overhead is only 0.24% of the 4MB LLC.

<sup>2</sup>Note that sharing ChargeCache across cores can result in even lower overheads. We leave the exploration of such designs to future work.



**Power Consumption.** ChargeCache is accessed on every *activate* and *precharge* command issued by the memory controller. On an *activate* command, ChargeCache is searched for the corresponding row address. On a *precharge* command, the address of the precharged row is inserted into ChargeCache. ChargeCache entries are periodically invalidated to ensure they do not exceed a specified *caching duration*. These three operations increase dynamic power consumption in the memory controller, and the ChargeCache storage increases static power consumption. Our analysis indicates that ChargeCache consumes 0.149 mW on average. This is only 0.23% of the average power consumption of the entire 4MB LLC. Note that we include the effect of this additional power consumption in our DRAM energy evaluations in Section 6.2. We conclude that ChargeCache incurs almost negligible chip area and power consumption overheads.

## 6.4 Sensitivity Studies

ChargeCache performance depends mainly on two variables: *HCRAC capacity* and *caching duration*. We observed that associativity has a negligible effect on ChargeCache performance. In our experiments, increasing the associativity of HCRAC from two to full-associativity improved the hit rate by only 2%. We analyze the hit rate and performance impact of *capacity* and *caching duration* in more detail.

### 6.4.1 ChargeCache capacity

Figure 6.3 shows the average hit rate versus capacity of ChargeCache for single-core and eight-core systems. The horizontal dashed lines indicate the maximum hit rate achievable with an unlimited-capacity ChargeCache. We observe that 128 entries is a sweet spot between hit rate and storage overhead. Such a configuration yields 38% and 66% hit rate for single-core and eight-core systems, respectively. The storage requirement for a 128-entry ChargeCache is only 672 bytes per core assuming our two-channel main memory (see Section 6.3).

Figure 6.4 shows the speedup with various ChargeCache capacities. Larger capacities provide higher performance thanks to the higher ChargeCache hit rate. However, they also incur higher hardware overhead. For a 128-entry capacity (672 bytes per-core), ChargeCache provides 8.8% performance improvement, and for a 1024-entry capacity (5376 bytes per-core) it provides 10.6% performance improvement. We conclude that ChargeCache is effective at various sizes, but its benefits start to diminish at higher capacities.

### 6.4.2 Caching duration

Increasing the *caching duration* may improve the hit rate by decreasing the number of invalidated entries. We evaluate several *caching durations* to determine the duration value that provides favorable performance. For each *caching duration*, Table 6.1 shows the *tRCD* and *tRAS* values which we obtain from our circuit-level SPICE simulations. We also provide the default timing parameters used as a baseline in the first row of the table.

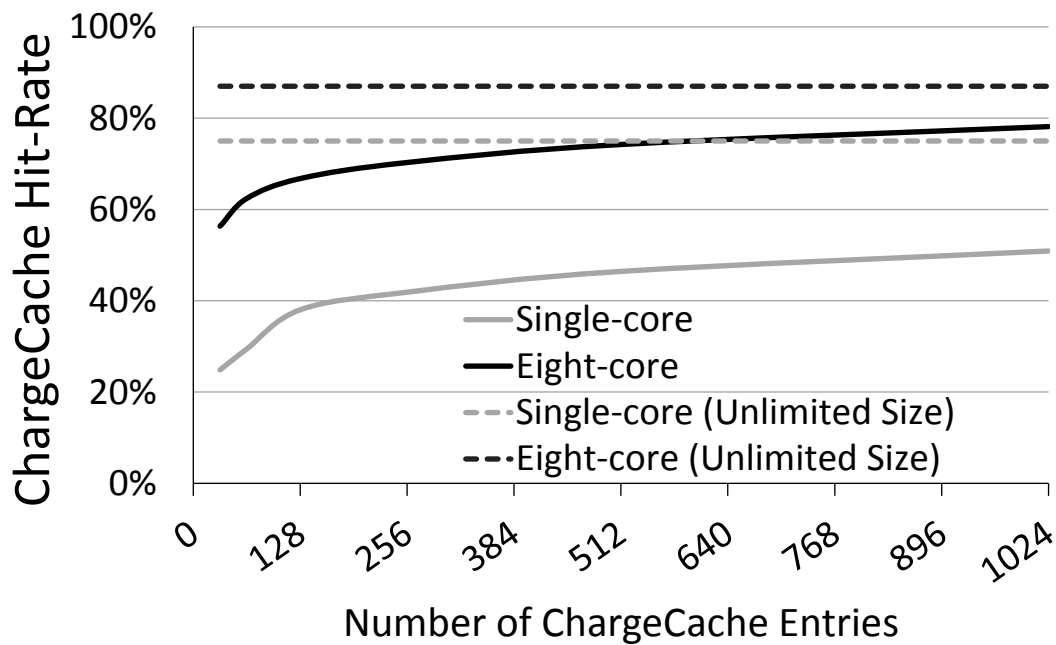


Figure 6.3: ChargeCache hit rate for single-core and eight-core systems at 1ms caching duration.

Table 6.1: tRCD and tRAS for different caching durations (determined via SPICE simulations)

Caching (ms)	Duration	tRCD (ns)	tRAS (ns)
N/A (Baseline)		13.75	35
1		8	22
4		9	24
16		11	28

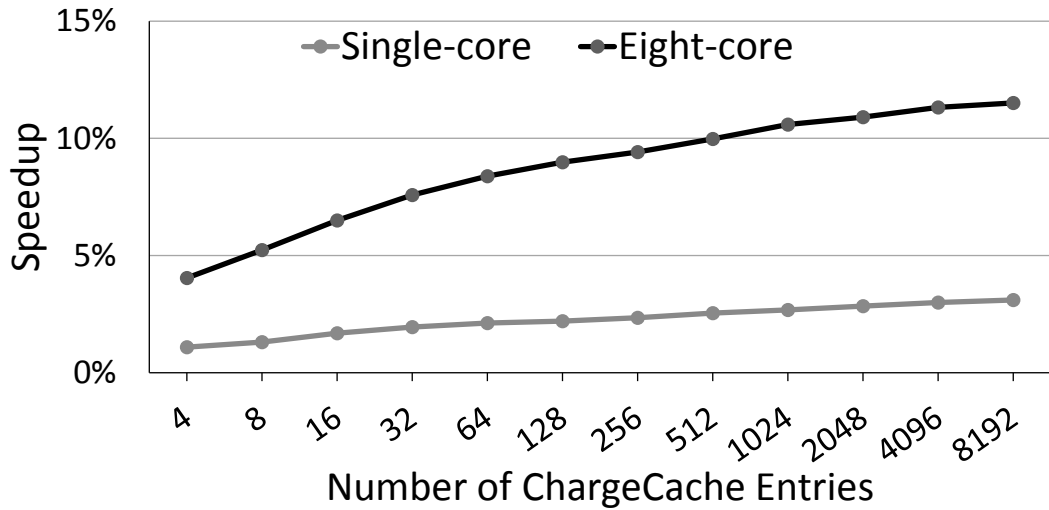


Figure 6.4: Speedup versus ChargeCache capacity.

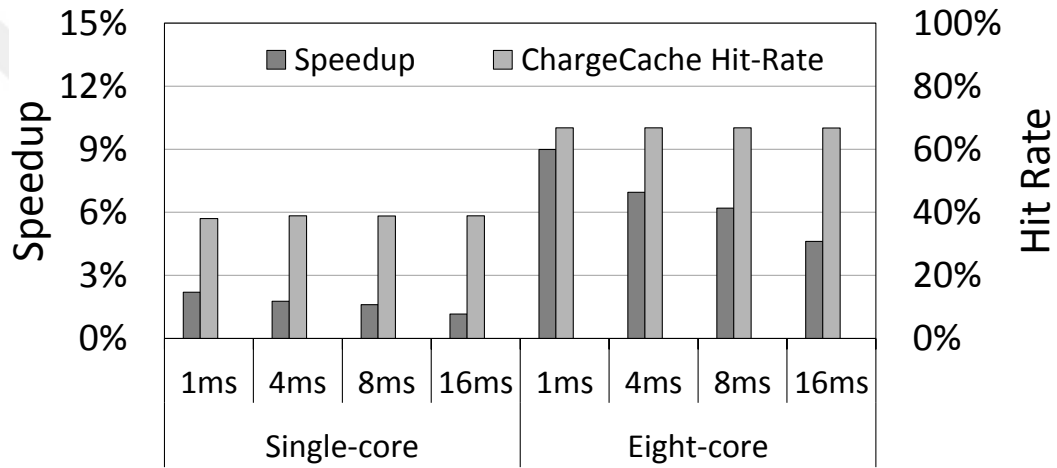


Figure 6.5: Speedup and ChargeCache hit rate for different *caching durations*

Figure 6.5 shows how ChargeCache speedup and ChargeCache hit rate vary with different *caching durations*. We make two observations. First, increasing the *caching duration* negatively affects the performance improvement of ChargeCache. This is because a longer caching duration leads to lower reductions in  $t_{RCD}$  and  $t_{RAS}$  (as Table 6.1 shows), thereby reducing the benefit of a ChargeCache hit. Second, ChargeCache hit rate increases slightly (by about 2%) for the single-core system but remains almost constant for the eight-core system when *caching duration* increases. The latter is due to the large number of bank conflicts in the 8-core system, as we explained in Section 3. With many bank conflicts, the aggregate number of precharge commands is high and ChargeCache evicts entries very frequently even with a 1ms *caching duration*. Thus, a longer *caching duration* does not have much effect on hit rate.

We conclude that, with a longer *caching duration*, the improvement in ChargeCache hit rate does not make up for the loss in the reduction of the timing parameters. We conclude that ChargeCache is effective for various *caching durations*, yet the empirically best *caching duration* is 1ms, which leads to the highest performance improvement.



## 7. DISCUSSION

### 7.1 Temperature Independence

Charge leakage rate of DRAM cells approximately doubles for every 10°C increase in the temperature [38, 47, 50, 57, 73]. This observation can be exploited to lower the DRAM latency when operating at low temperatures. A previous study, Adaptive Latency DRAM (AL-DRAM) [47], proposes a mechanism to improve system performance by reducing the DRAM timing parameters at low operating temperature. It is based on the premise that DRAM typically does not operate at temperatures close to the worst-case temperature (85° C) even when it is heavily accessed. However, new 3D-stacked DRAM technologies such as HMC [30], HBM [32], WideIO [17] may operate at significantly higher temperatures due to tight integration of multiple stack layers [4, 46, 71]. Therefore, *dynamic latency scaling* techniques such as AL-DRAM may be less useful in these scenarios.

ChargeCache is *not* based on the charge difference that occurs due to temperature dependence. Rather, we exploit the high level of charge in recently-precharged rows to reduce timing parameters during later accesses to such rows. After conducting tests to determine the reduction in timing parameters (for ChargeCache hits) at *worst-case* temperatures, we find that these timing parameters can be reduced independently of the operating temperature. Dynamic latency scaling can still be used in conjunction with ChargeCache at low temperatures to reduce the access latency even further.

### 7.2 Applicability to Other DRAM Standards

Although we evaluate only DDR3-based main memory within the scope of this thesis, implementing ChargeCache for other DRAM standards is straightforward. In theory, ChargeCache is applicable to any memory technology where cells are volatile (leak charge over time). However, the memory interface can prevent the implementation of ChargeCache entirely in the memory controller. For example, RL-DRAM [101] is a DRAM type incompatible with ChargeCache. In RL-DRAM, read and write operations are directly handled by *READ* and *WRITE* commands without explicitly activating and precharging DRAM rows. Hence, the RL-DRAM memory controller does not have control over the activation delay of the rows and the timing parameters *tRCD* and *tRAS* do not exist.

However, ChargeCache can be used with to a large set of specifications derived from DDR (DDR, GDDR, LPDDR, etc.) in a manner similar to the mechanism described in this work, without modifying the DRAM architecture at all. All of these memories require *ACT* and *PRE* commands to explicitly open and close DRAM rows. Using ChargeCache with 3D-stacked memories [46, 52] such as WideIO [17], HBM [32] and HMC [30] is also straightforward. The difference is that the DRAM controller,

and hence ChargeCache, may be implemented in the logic layer of the 3D-stacked memory chip instead of the processor chip.



## 8. RELATED WORK

To our knowledge, this work is the first to (i) show that applications typically exhibit significant *Row-level Temporal Locality (RLTL)* and (ii) exploit this locality to improve system performance by reducing the latency of requests to recently-accessed rows.

We have already qualitatively and quantitatively (in Sections 3 and 6) compared ChargeCache to NUAT [82], which reduces access latency to *only* recently-refreshed rows. We have shown that ChargeCache can provide significantly higher average latency reduction than NUAT because RLTL is usually high, whereas the fraction of accesses to rows that are recently-refreshed is typically low.

Other previous works have proposed techniques to reduce performance degradation caused by long DRAM latencies. They focused on 1) enhancing the DRAM, 2) exploiting variations in manufacturing process and operating conditions, 3) developing several memory scheduling policies. We briefly summarize how ChargeCache differs from these works.

**Enhancing DRAM Architecture.** Lee et al. propose Tiered-Latency DRAM (TL-DRAM) [48] which divides each subarray into near and far segments using isolation transistors. With TL-DRAM, the memory controller accesses the near segment with lower latency since the isolation transistor reduces bitline capacitance in that segment. Our mechanism could be implemented on top of TL-DRAM to reduce the access latency for both the near and far segment. Kim et al. unlock parallelism among subarrays at low cost with SALP [41]. The goal of SALP is to reduce DRAM latency by providing more parallelism to reduce the impact of bank conflicts. O et al [67] propose a DRAM architecture where sense amplifiers are decoupled from bitlines to mitigate precharge latency. Choi et al [12] propose to utilize multiple DRAM cells to store a single bit when sufficient DRAM capacity is available. By using multiple cells, they reduce activation, precharge and refresh latencies. Other works [8, 9, 27, 78–80, 85, 96] also propose new DRAM architectures to lower DRAM latency.

Unlike ChargeCache, all these works require changes to the DRAM architecture itself. The approaches taken by these works are largely orthogonal and ChargeCache could be implemented together with any of these mechanisms to further improve the DRAM latency.

**Exploiting Process and Operating Condition Variations.** Recent studies [6, 47] proposed methods to reduce the safety margins of the DRAM timing parameters when operating conditions are appropriate (i.e., not worst-case). Unlike these works, ChargeCache is largely independent of operating conditions like temperature, as discussed in Section 7.1, and is orthogonal to these latency reduction mechanisms.

**Memory Request Scheduling Policies.** Memory request scheduling policies

(e.g., [39, 40, 44, 62, 63, 76, 87–89, 93]) reduce the average DRAM access latency by improving DRAM parallelism, row-buffer locality and fairness in especially multi-core systems. ChargeCache can be employed in conjunction with the scheduling policy that best suits the application and the underlying architecture.





## 9. CONCLUSION

We introduce ChargeCache, a new, low-overhead mechanism that dynamically reduces the DRAM timing parameters for recently-accessed DRAM rows. ChargeCache exploits two key observations that we demonstrate in this work: 1) a recently-accessed DRAM row has cells with high amount of charge and thus can be accessed faster, 2) many applications repeatedly access rows that are recently-accessed.

Our extensive evaluations of ChargeCache on both single-core and multi-core systems show that it provides significant performance benefit and DRAM energy reduction at very modest hardware overhead. ChargeCache requires no modifications to the existing DRAM chips and occupies only a small area on the memory controller.

We conclude that ChargeCache is a simple yet efficient mechanism to dynamically reduce DRAM latency, which significantly improves both the performance and energy efficiency of modern systems.



## REFERENCES

- [1] **Awasthi, M., Nellans, D. W., Balasubramonian, R., and Davis, A.** Prediction based DRAM row-buffer management in the many-core era. In *PACT* (2011).
- [2] **Baer, J.-L., and Chen, T.-F.** An effective on-chip preloading scheme to reduce data access penalty. In *ICS* (1991).
- [3] **Bekerman, M., Jourdan, S., Ronen, R., Kirshenboim, G., Rappoport, L., Yoaz, A., and Weiser, U.** Correlated load-address predictors. In *ISCA* (1999).
- [4] **Black, B., Annavaram, M., Brekelbaum, N., DeVale, J., Jiang, L., Loh, G. H., McCauley, D., Morrow, P., Nelson, D. W., Pantuso, D., Reed, P., Rupley, J., Shankar, S., Shen, J., and Webb, C.** Die stacking (3D) microarchitecture. In *MICRO* (2006).
- [5] **Chandra, D., Guo, F., Kim, S., and Solihin, Y.** Predicting inter-thread cache contention on a chip multi-processor architecture. In *HPCA* (2005).
- [6] **Chandrasekar, K., Goossens, S., Weis, C., Koedam, M., Akesson, B., Wehn, N., and Goossens, K.** Exploiting expendable process-margins in DRAMs for run-time performance optimization. In *DATE* (2014).
- [7] **Chandrasekar, K., Weis, C., Akesson, B., Wehn, N., and Goossens, K.** Towards variation-aware system-level power estimation of DRAMs: an empirical approach. In *DAC* (2013).
- [8] **Chang, K. K.-W., et al.** Low-cost inter-linked subarrays (LISA): Enabling fast inter-subarray data movement in DRAM. In *HPCA* (2016).
- [9] **Chang, K. K.-W., Lee, D., Chishti, Z., Alameldeen, A. R., Wilkerson, C., Kim, Y., and Mutlu, O.** Improving DRAM performance by parallelizing refreshes with accesses. In *HPCA* (2014).
- [10] **Charney, M. J., and Puzak, T. R.** Prefetching and memory system behavior of the SPEC95 benchmark suite. *IBM JRD* (1997).
- [11] **Chatterjee, N., Shevgoor, M., Balasubramonian, R., Davis, A., Fang, Z., Illikkal, R., and Iyer, R.** Leveraging heterogeneity in DRAM main memories to accelerate critical word access. In *MICRO* (2012).

- [12] **Choi, J., Shin, W., Jang, J., Suh, J., Kwon, Y., Moon, Y., and Kim, L.-S.** Multiple clone row DRAM: a low latency and area optimized DRAM. In *ISCA* (2015).
- [13] **Chou, Y., Fahs, B., and Abraham, S.** Microarchitecture optimizations for exploiting memory-level parallelism. In *ISCA* (2004).
- [14] **Das, R., Ausavarungnirun, R., Mutlu, O., Kumar, A., and Azimi, M.** Application-to-core mapping policies to reduce memory system interference in multi-core systems. In *HPCA* (2013).
- [15] **Ding, W., Guttman, D., and Kandemir, M.** Compiler support for optimizing memory bank-level parallelism. In *MICRO* (2014).
- [16] **Duong, N., Zhao, D., Kim, T., Cammarota, R., Valero, M., and Veidenbaum, A. V.** Improving cache management policies using dynamic reuse distances. In *MICRO* (2012).
- [17] **Dutoit, D., Bernard, C., Cheramy, S., Clermidy, F., Thonnart, Y., Vivet, P., Freund, C., Guerin, V., Guilhot, S., Lecomte, S., et al.** A 0.9 pJ/bit, 12.8 GByte/s WideIO memory interface in a 3D-IC NoC-based MPSoC. In *VLSIT* (2013).
- [18] **Ebrahimi, E., Lee, C. J., Mutlu, O., and Patt, Y. N.** Prefetch-aware shared resource management for multi-core systems. In *ISCA* (2011).
- [19] **Ebrahimi, E., Miftakhutdinov, R., Fallin, C., Lee, C. J., Joao, J. A., Mutlu, O., and Patt, Y. N.** Parallel application memory scheduling. In *MICRO* (2011).
- [20] **Ebrahimi, E., Mutlu, O., Lee, C. J., and Patt, Y. N.** Coordinated control of multiple prefetchers in multi-core systems. In *MICRO* (2009).
- [21] **Ebrahimi, E., Mutlu, O., and Patt, Y. N.** Techniques for bandwidth-efficient prefetching of linked data structures in hybrid prefetching systems. In *HPCA* (2009).
- [22] **Eickemeyer, R. J., and Vassiliadis, S.** A load-instruction unit for pipelined processors. *IBM JRD* (1993).
- [23] **Eyerman, S., and Eeckhout, L.** System-level performance metrics for multiprogram workloads. *IEEE Micro* (2008).
- [24] **Ghosh, M., and Lee, H.-H. S.** Smart refresh: An enhanced memory controller design for reducing energy in conventional and 3D die-stacked DRAMs. In *MICRO* (2007).
- [25] **Glew, A.** MLP yes! ILP no. In *ASPLOS WACI* (1998).
- [26] **GSI.** Low latency DRAMs. <http://www.gsitechnology.com>.
- [27] **Gulur, N. D., Manikantan, R., Mehendale, M., and Govindarajan, R.** Multiple sub-row buffers in DRAM: unlocking performance and energy improvement opportunities. In *ICS* (2012).

- [28] **Hamerly, G., Perelman, E., Lau, J., and Calder, B.** Simpoint 3.0: Faster and more flexible program phase analysis. *JILP* (2005).
- [29] **Hassan, H., Pekhimenko, G., Vijaykumar, N., Seshadri, V., Lee, D., Ergin, O., and Mutlu, O.** Chargecache: Reducing dram latency by exploiting row access locality. In *2016 IEEE International Symposium on High Performance Computer Architecture (HPCA)* (2016), IEEE, pp. 581–593.
- [30] **Hybrid Memory Cube Consortium.** Hybrid memory cube specification 2.0. Tech. rep., November 2014.
- [31] **Iyer, R., Zhao, L., Guo, F., Illikkal, R., Makineni, S., Newell, D., Solihin, Y., Hsu, L., and Reinhardt, S.** QoS policies and architecture for cache/memory in CMP platforms. In *SIGMETRICS* (2007).
- [32] **JEDEC.** High bandwidth memory (HBM) DRAM.
- [33] **Jeong, M. K., Yoon, D. H., Sunwoo, D., Sullivan, M., Lee, I., and Erez, M.** Balancing DRAM locality and parallelism in shared memory CMP systems. In *HPCA* (2012).
- [34] **JESD209, J. S.** Low power double data rate (lpddr) sdram specification. *JEDEC Solid State Technology Association 8* (2007).
- [35] **Jevdjic, D., Loh, G. H., Kaynak, C., and Falsafi, B.** Unison cache: A scalable and effective die-stacked DRAM cache. In *MICRO* (2014).
- [36] **Joao, J. A., Suleman, M. A., Mutlu, O., and Patt, Y. N.** Bottleneck identification and scheduling in multithreaded applications. In *ASPLOS* (2012).
- [37] **Kandemir, M., Zhao, H., Tang, X., and Karakoy, M.** Memory row reuse distance and its role in optimizing application performance. In *SIGMETRICS* (2015).
- [38] **Khan, S., Lee, D., Kim, Y., Alameldeen, A. R., Wilkerson, C., and Mutlu, O.** The efficacy of error mitigation techniques for DRAM retention failures: a comparative experimental study. In *SIGMETRICS* (2014).
- [39] **Kim, Y., Han, D., Mutlu, O., and Harchol-Balter, M.** ATLAS: A scalable and high-performance scheduling algorithm for multiple memory controllers. In *HPCA* (2010).
- [40] **Kim, Y., Papamichael, M., Mutlu, O., and Harchol-Balter, M.** Thread cluster memory scheduling: Exploiting differences in memory access behavior. In *MICRO* (2010).
- [41] **Kim, Y., Seshadri, V., Lee, D., Liu, J., and Mutlu, O.** A case for exploiting subarray-level parallelism (SALP) in DRAM. In *ISCA* (2012).
- [42] **Kim, Y., Yang, W., and Mutlu, O.** Ramulator: A fast and extensible DRAM simulator. In *CAL* (2015).

- [43] **Lee, C. J., Mutlu, O., Narasiman, V., and Patt, Y. N.** Prefetch-aware DRAM controllers. In *MICRO* (2008).
- [44] **Lee, C. J., Narasiman, V., Ebrahimi, E., Mutlu, O., and Patt, Y. N.** DRAM-aware last-level cache writeback: Reducing write-caused interference in memory systems. *UT-Austin, HPS, Tech. Report* (2010).
- [45] **Lee, C. J., Narasiman, V., Mutlu, O., and Patt, Y. N.** Improving memory bank-level parallelism in the presence of prefetching. In *MICRO* (2009).
- [46] **Lee, D., Ghose, S., Pekhimenko, G., Khan, S., and Mutlu, O.** Simultaneous multi-layer access: Improving 3D-stacked memory bandwidth at low cost. *TACO* (2016).
- [47] **Lee, D., Kim, Y., Pekhimenko, G., Khan, S., Seshadri, V., Chang, K., and Mutlu, O.** Adaptive-latency DRAM: Optimizing DRAM timing for the common-case. In *HPCA* (2015).
- [48] **Lee, D., Kim, Y., Seshadri, V., Liu, J., Subramanian, L., and Mutlu, O.** Tiered-latency DRAM: A low latency and low cost DRAM architecture. In *HPCA* (2013).
- [49] **Li, S., Ahn, J. H., Strong, R. D., Brockman, J. B., Tullsen, D. M., and Jouppi, N. P.** McPAT: an integrated power, area, and timing modeling framework for multicore and manycore architectures. In *MICRO* (2009).
- [50] **Liu, J., Jaiyen, B., Kim, Y., Wilkerson, C., and Mutlu, O.** An experimental study of data retention behavior in modern DRAM devices: Implications for retention time profiling mechanisms. In *ISCA* (2013).
- [51] **Liu, J., Jaiyen, B., Veras, R., and Mutlu, O.** RAIDR: Retention-aware intelligent DRAM refresh. In *ISCA* (2012).
- [52] **Loh, G. H.** 3D-stacked memory architectures for multi-core processors. In *ISCA* (2008).
- [53] **Lotfi-Kamran, P., Grot, B., Ferdman, M., Volos, S., Kocberber, O., Picorel, J., Adileh, A., Jevdjic, D., Idgunji, S., Ozer, E., et al.** Scale-out processors. In *ISCA* (2012).
- [54] **Luk, C.-K., Cohn, R., Muth, R., Patil, H., Klauser, A., Lowney, G., Wallace, S., Reddi, V. J., and Hazelwood, K.** Pin: building customized program analysis tools with dynamic instrumentation. In *PLDI* (2005).
- [55] **Matick, R. E., and Schuster, S. E.** Logic-based eDRAM: Origins and rationale for use. *IBM Journal of Research and Development* 49, 1 (2005), 145.
- [56] **Micron Technology.** *4Gb: x4, x8, x16 DDR3 SDRAM*, 2011.

- [57] **Mori, Y., Ohyu, K., Okonogi, K., and Yamada, R.-i.** The origin of variable retention time in DRAM. In *IEDM* (2005).
- [58] **Moscibroda, T., and Mutlu, O.** Memory performance attacks: Denial of memory service in multi-core systems. In *USENIX Security* (2007).
- [59] **Muralidhara, S. P., Subramanian, L., Mutlu, O., Kandemir, M., and Moscibroda, T.** Reducing memory interference in multicore systems via application-aware memory channel partitioning. In *MICRO* (2011).
- [60] **Mutlu, O., Kim, H., and Patt, Y. N.** Address-value delta (AVD) prediction: Increasing the effectiveness of runahead execution by exploiting regular memory allocation patterns. In *MICRO* (2005).
- [61] **Mutlu, O., Kim, H., and Patt, Y. N.** Techniques for efficient processing in runahead execution engines. In *ISCA* (2005).
- [62] **Mutlu, O., and Moscibroda, T.** Stall-time fair memory access scheduling for chip multiprocessors. In *MICRO* (2007).
- [63] **Mutlu, O., and Moscibroda, T.** Parallelism-aware batch scheduling: Enhancing both performance and fairness of shared DRAM systems. In *ISCA* (2008).
- [64] **Mutlu, O., Stark, J., Wilkerson, C., and Patt, Y. N.** Runahead execution: An alternative to very large instruction windows for out-of-order processors. In *HPCA* (2003).
- [65] **Mutlu, O., and Subramanian, L.** Research problems and opportunities in memory systems. *SUPERFRI* (2014).
- [66] **Nair, P., Chou, C.-C., and Qureshi, M. K.** A case for refresh pausing in DRAM memory systems. In *HPCA* (2013).
- [67] **O, S., Son, Y. H., Kim, N. S., and Ahn, J. H.** Row-buffer decoupling: a case for low-latency DRAM microarchitecture. In *ISCA* (2014).
- [68] **Pai, V. S., and Adve, S.** Code transformations to improve memory parallelism. In *MICRO* (1999).
- [69] **Palacharla, S., and Kessler, R.** Evaluating stream buffers as a secondary cache replacement. In *ISCA* (1994).
- [70] **Patt, Y. N., Hwu, W.-m., and Shebanow, M.** HPS, a new microarchitecture: rationale and introduction. In *MICRO* (1985).
- [71] **Puttaswamy, K., and Loh, G. H.** Thermal analysis of a 3d die-stacked high-performance microprocessor. In *GLSVLSI* (2006).
- [72] **Qureshi, M. K., Jaleel, A., Patt, Y. N., Steely, S. C., and Emer, J.** Adaptive insertion policies for high performance caching. In *ISCA* (2007).

- [73] **Qureshi, M. K., Kim, D.-H., Khan, S., Nair, P. J., and Mutlu, O.** AVATAR: A variable-retention-time (VRT) aware refresh for DRAM systems. In *DSN* (2015).
- [74] **Qureshi, M. K., Suleman, M. A., and Patt, Y. N.** Line distillation: Increasing cache capacity by filtering unused words in cache lines. In *HPCA* (2007).
- [75] **Rau, B. R.** Pseudo-randomly interleaved memory. In *ISCA* (1991).
- [76] **Rixner, S., Owens, J. D., Mattson, P., Kapasi, U. J., and Dally, W. J.** Memory access scheduling. In *ISCA* (2000).
- [77] **Sato, Y., Suzuki, T., Aikawa, T., Fujioka, S., Fujieda, W., Kobayashi, H., Ikeda, H., Nagasawa, T., Funyu, A., Fuji, Y., et al.** Fast cycle RAM (FCRAM); a 20-ns random row access, pipe-lined operating DRAM. In *VLSI Circuits* (1998).
- [78] **Seshadri, V., Hsieh, K., Boroumand, A., Lee, D., et al.** Fast bulk bitwise AND and OR in DRAM. In *CAL* (2015).
- [79] **Seshadri, V., Kim, Y., Fallin, C., Lee, D., Ausavarungnirun, R., Pekhimenko, G., Luo, Y., Mutlu, O., Gibbons, P. B., Kozuch, M. A., et al.** RowClone: Fast and energy-efficient in-DRAM bulk data copy and initialization. In *MICRO* (2013).
- [80] **Seshadri, V., Mullins, T., Boroumand, A., Mutlu, O., Gibbons, P. B., Kozuch, M. A., and Mowry, T. C.** Gather-scatter DRAM: In-DRAM address translation to improve the spatial locality of non-unit strided accesses. In *MICRO* (2015).
- [81] **Seshadri, V., Mutlu, O., Kozuch, M. A., and Mowry, T. C.** The evicted-address filter: A unified mechanism to address both cache pollution and thrashing. In *PACT* (2012).
- [82] **Shin, W., Yang, J., Choi, J., , and Kim, L.-S.** NUAT: A non-uniform access time memory controller. In *HPCA* (2014).
- [83] **Smith, B. J.** A pipelined, shared resource MIMD computer. In *ICPP* (1978).
- [84] **Snavely, A., and Tullsen, D. M.** Symbiotic jobscheduling for a simultaneous multithreading processor. *ASPLOS* (2000).
- [85] **Son, Y. H., Seongil, O., Ro, Y., Lee, J. W., and Ahn, J. H.** Reducing memory access latency with asymmetric DRAM bank organizations. In *ISCA* (2013).
- [86] **Srinath, S., Mutlu, O., Kim, H., and Patt, Y. N.** Feedback directed prefetching: Improving the performance and bandwidth-efficiency of hardware prefetchers. In *HPCA* (2007).
- [87] **Subramanian, L., Lee, D., Seshadri, V., Rastogi, H., and Mutlu, O.** The blacklisting memory scheduler: Achieving high performance and fairness at low cost. In *ICCD* (2014).



- [88] **Subramanian, L., Seshadri, V., Ghosh, A., Khan, S., and Mutlu, O.** The application slowdown model: Quantifying and controlling the impact of inter-application interference at shared caches and main memory. In *MICRO* (2015).
- [89] **Subramanian, L., Seshadri, V., Kim, Y., Jaiyen, B., and Mutlu, O.** MISE: Providing performance predictability and improving fairness in shared main memory systems. In *HPCA* (2013).
- [90] **Suleman, M. A., Mutlu, O., Qureshi, M. K., and Patt, Y. N.** Accelerating critical section execution with asymmetric multi-core architectures. In *ASPLOS* (2009).
- [91] **Thornton, J. E.** Parallel operation in the Control Data 6600. In *Fall Joint Computer Conference* (1964).
- [92] **Tomasulo, R. M.** An efficient algorithm for exploiting multiple arithmetic units. *IBM JRD* (1967).
- [93] **Usui, H., Subramanian, L., Chang, K. K.-W., and Mutlu, O.** DASH: Deadline-aware high-performance memory scheduler for heterogeneous systems with hardware accelerators. *TACO* (2016).
- [94] **Wilkes, M. V.** The memory gap and the future of high performance memories. In *ACM SIGARCH Computer Architecture News* 29 (2001).
- [95] **Wulf, W. A., and McKee, S. A.** Hitting the memory wall: implications of the obvious. In *ACM SIGARCH Computer Architecture News* 23 (1995).
- [96] **Zhang, T., Chen, K., Xu, C., Sun, G., Wang, T., and Xie, Y.** Half-DRAM: a high-bandwidth and low-power DRAM architecture from the rethinking of fine-grained activation. In *ISCA* (2014).
- [97] **Zhang, Z., Zhu, Z., and Zhang, X.** A permutation-based page interleaving scheme to reduce row-buffer conflicts and exploit data locality. In *MICRO* (2000).
- [98] **Zhao, W., and Cao, Y.** New generation of predictive technology model for sub-45 nm early design exploration. *IEEE TED* (2006).
- [99] **Zuravleff, W. K., and Robinson, T.** Controller for a synchronous dram that maximizes throughput by allowing memory requests and commands to be issued out of order, 1997. US Patent 5,630,096.
- [100] Joint electron device engineering council (JEDEC). <https://www.jedec.org>. Accessed: 2016-08-02.
- [101] Micron RLDRAM 2 and 3 specifications. <http://www.micron.com/products/dram/rldram-memory>.
- [102] Predictive technology model (ptm). <http://ptm.asu.edu/>.
- [103] Rambus DRAM power model (2010). <http://www.rambus.com/energy>.

[104] Ramulator (source code). <https://github.com/CMU-SAFARI/ramulator>.

[105] SPEC CPU 2006. <http://www.spec.org/cpu2006>.

[106] STREAM Benchmark. <http://www.streambench.org/>.

[107] TPC. <http://www.tpc.org/>.



## CURRICULUM VITAE

**Name-Surname** : Hasan Hassan  
**Nationality** : Turkish, Bulgarian  
**Birth Date and Place** : July 4, 1991, Kubrat (Bulgaria)  
**e-mail** : hasanibrahimhasan@gmail.com

### EDUCATION:

- **Bachelor's Degree:** 2014, TOBB University of Economics and Technology,  
Engineering Faculty, Computer Engineering

### EXPERIENCE:

Year	Place	Position
2011	Kasirga Microprocessors Research Laboratory	Computer Engineering Intern
2012	Taleworlds Entertainment	Computer Engineering Intern
2013	Berlin Technical University	Computer Engineering Intern
2014-2016	TOBB University of Economics and Technology	Distinguished Scholar
2015	Carnegie Mellon University	Research Scholar

**LANGUAGES:** Turkish, English, Bulgarian

### RELATED PUBLICATIONS:

- Hassan, H., Pekhimenko, G., Vijaykumar, N., Seshadri, V., Lee, D., Ergin, O., and Mutlu, O. (2016, March). ChargeCache: Reducing DRAM latency by exploiting row access locality. In 2016 IEEE International Symposium on High Performance Computer Architecture (HPCA) (pp. 581-593). IEEE.
- Chang, K. K., Kashyap, A., Hassan, H., Ghose, S., Hsieh, K., Lee, D., ... and Mutlu, O. (2016, June). Understanding Latency Variation in Modern DRAM Chips: Experimental Characterization, Analysis, and Optimization. In Proceedings of the 2016 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science (pp. 323-336). ACM.