

In-DRAM Bulk Bitwise Execution Engine

Vivek Seshadri
 Microsoft Research India
 visesha@microsoft.com

Onur Mutlu
 ETH Zürich
 onur.mutlu@inf.ethz.ch

Abstract

Many applications heavily use bitwise operations on large bitvectors as part of their computation. In existing systems, performing such *bulk bitwise operations* requires the processor to transfer a large amount of data on the memory channel, thereby consuming high latency, memory bandwidth, and energy. In this paper, we describe *Ambit*, a recently-proposed mechanism to perform bulk bitwise operations *completely* inside main memory. *Ambit* exploits the internal organization and analog operation of DRAM-based memory to achieve low cost, high performance, and low energy. *Ambit* exposes a new bulk bitwise execution model to the host processor. Evaluations show that *Ambit* significantly improves the performance of several applications that use bulk bitwise operations, including databases.

Index Terms: Processing using Memory, DRAM, Bulk Copy, Bulk Initialization, Bulk Bitwise Operations, Performance, Energy Efficiency

1. Introduction

Many applications trigger *bulk bitwise operations*, i.e., bitwise operations on large bit vectors [84, 146]. In databases, bitmap indices [28, 113], which heavily use bulk bitwise operations, are more efficient than B-trees for many queries [28, 5, 149]. In fact, many real-world databases [114, 10, 5, 11] support bitmap indices. A recent work, *WideTable* [98], designs an entire database around a technique called *BitWeaving* [97], which accelerates *scans* completely using bulk bitwise operations. Microsoft recently open-sourced a technology called *BitFunnel* [47] that accelerates the document filtering portion of web search. *BitFunnel* relies on fast bulk bitwise AND operations. Bulk bitwise operations are also prevalent in DNA sequence alignment [21, 150, 18, 77, 111, 129, 17], encryption algorithms [143, 51, 106], graph processing [96, 15, 112], networking [146], and machine learning [38]. Thus, accelerating bulk bitwise operations can significantly boost the performance of various important applications.

In existing systems, a bulk bitwise operation requires a large amount of data to be transferred on the memory channel. Such large data transfers result in high latency, bandwidth, and energy consumption. In fact, our experiments on a multi-core Intel Skylake [1] and an

NVIDIA GeForce GTX 745 [6] show that the available memory bandwidth of these systems limits the throughput of bulk bitwise operations. Recent works (e.g., [15, 14, 154, 41, 49, 59, 58, 25, 77, 24, 102, 119]) propose processing in the logic layer of 3D-stacked DRAM, which stacks DRAM layers on top of a logic layer (e.g., Hybrid Memory Cube [63, 8], High Bandwidth Memory [7, 88]). While the logic layer in 3D-stacked memory has much higher bandwidth than traditional systems, it still cannot exploit the maximum internal bandwidth available inside a DRAM chip [88] (Section 7).

In this paper, we describe Ambit [134], a new mechanism proposed by recent work that performs bulk bitwise operation completely inside main memory. Ambit is an instance of the recently-introduced notion called *Processing using Memory* [136]. In contrast to Processing in Memory architectures [15, 14, 154, 49, 157, 141, 58, 43, 108, 57, 42, 41, 20, 16, 88, 139, 137, 118, 85, 115, 46, 39, 132, 133, 59, 53, 52, 24, 102, 25, 112, 26] that add extra computational logic closer to main memory, the idea behind Processing using Memory is to exploit the existing structure and organization of memory devices with minimal changes to provide additional functionality.

Along these lines, Ambit uses the analog operation principles of DRAM technology to perform bulk bitwise operations *completely inside the memory array*. With modest changes to the DRAM design, Ambit can exploit 1) the maximum internal bandwidth available *inside* each DRAM array, and 2) the memory-level parallelism [109, 82, 45, 81, 110, 86] *across* multiple DRAM arrays to enable one to two orders of magnitude improvement in raw throughput and energy consumption of bulk bitwise operations. Ambit exposes a *Bulk Bitwise Execution Model* to the host processor. We show that real-world applications that heavily use bulk bitwise operations can use this model to achieve significant improvements in performance and energy efficiency.

In this paper, we discuss the following main concepts.

- As Ambit builds on top of modern DRAM architecture, we first provide a brief background on modern DRAM organization and operation that is sufficient to understand the mechanisms proposed by Ambit (Section 2).
- We describe the different components of Ambit, their design and implementation, and execution models to expose Ambit to the host system in detail (Sections 3–5).
- We describe quantitative evaluations showing that 1) Ambit works reliably even under significant process variation, and 2) improves performance and energy-efficiency compared to existing systems for real workloads (Sections 6–8).

2. Background on DRAM

In this section, we describe the necessary background to understand modern DRAM architecture and its implementation. This paper builds on our previous book chapter that introduces the notion of *Processing using Memory* [136]. Since that chapter provides a detailed background on DRAM, this section is mostly reproduced from the DRAM background section from chapter. While we focus our attention primarily on commodity DRAM design

(i.e., the DDRx interface), most DRAM architectures use very similar design approaches and vary only in higher-level design choices [83]. As a result, Ambit can be extended to any DRAM architecture. There has been significant recent research in DRAM architectures and the interested reader can find details about various aspects of DRAM in multiple recent publications [82, 92, 91, 135, 33, 83, 100, 73, 30, 71, 101, 55, 121, 90, 117, 56, 31, 88, 130, 79, 87, 29, 133, 93, 72, 70, 54, 145, 76, 75, 74, 116].

At the end of this section, we provide a brief overview of RowClone [133], a prior work that enables the memory controller to perform row-wide copy and initialization operations completely inside DRAM. Ambit exploits RowClone to reduce the overhead of some of its bulk data copy and initialization operations.

2.1. High-level Organization of the Memory System

Figure 1 shows the organization of the memory subsystem in a modern computing system. At a high level, each processor chip is connected to one or more off-chip memory *channels*. Each memory channel consists of its own set of *command*, *address*, and *data* buses. Depending on the design of the processor, there can be either an independent memory controller for each memory channel or a single memory controller for all memory channels. All memory modules connected to a channel share the buses of the channel. Each memory module consists of many DRAM devices (or chips). Most of this section is dedicated to describing the design of a modern DRAM chip. In Section 2.3, we present more details of the module organization of commodity DRAM.

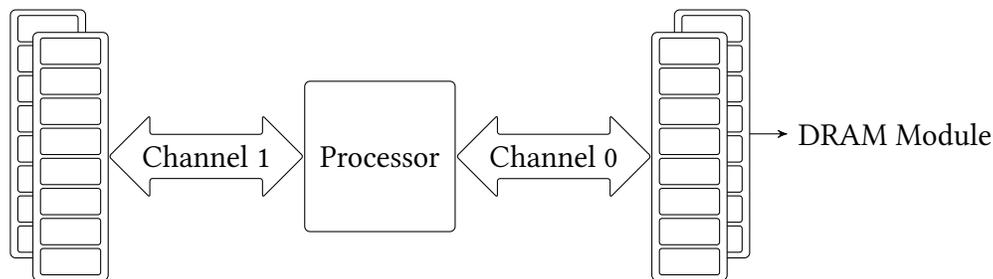


Figure 1: High-level organization of the memory subsystem

2.2. DRAM Chip

A modern DRAM chip consists of a hierarchy of structures: DRAM *cells*, *tiles/MATs*, *subarrays*, and *banks*. In this section, we describe the design of a modern DRAM chip in a bottom-up fashion, starting from a single DRAM cell and its operation.

2.2.1. DRAM Cell and Sense Amplifier

At the lowest level, DRAM technology uses capacitors to store information. Specifically, it uses the two extreme states of a capacitor, namely, the *empty* and the *fully charged* states to store a single bit of information. For instance, an empty capacitor can denote a logical

value of 0, and a fully charged capacitor can denote a logical value of 1. Figure 2 shows the two extreme states of a capacitor.

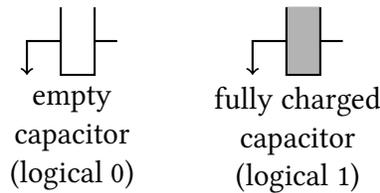


Figure 2: Two states of a capacitor

Unfortunately, the capacitors used for DRAM chips are small, and will get smaller with each new generation. As a result, the amount of charge that can be stored in the capacitor, and hence the difference between the two states is also very small. In addition, the capacitor can potentially lose its state after it is accessed. Therefore, to extract the state of the capacitor, DRAM manufacturers use a component called *sense amplifier*.

Figure 3 shows a sense amplifier. A sense amplifier contains two inverters which are connected together such that the output of one inverter is connected to the input of the other and vice versa. The sense amplifier also has an enable signal that determines if the inverters are active. When enabled, the sense amplifier has two stable states, as shown in Figure 4. In both these stable states, each inverter takes a logical value and feeds the other inverter with the negated input.

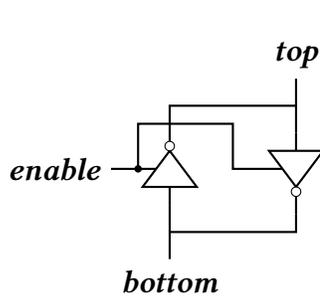


Figure 3: Sense amplifier

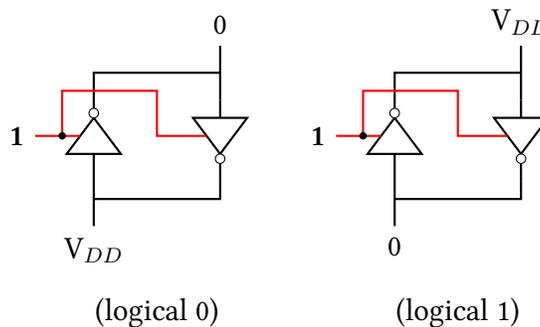


Figure 4: Stable states of a sense amplifier

Figure 5 shows the operation of the sense amplifier, starting from a disabled state. In the initial disabled state, we assume that the voltage level of the top terminal (V_a) is higher than that of the bottom terminal (V_b). When the sense amplifier is enabled in this state, it *senses* the difference between the two terminals and *amplifies* the difference until it reaches one of the stable states (hence the name “sense amplifier”).

2.2.2. DRAM Cell Operation: The ACTIVATE-PRECHARGE cycle

DRAM technology uses a simple mechanism that converts the logical state of a capacitor into a logical state of the sense amplifier. Data can then be accessed from the sense amplifier (since it is in a stable state). Figure 6 shows 1) the connection between a DRAM cell and

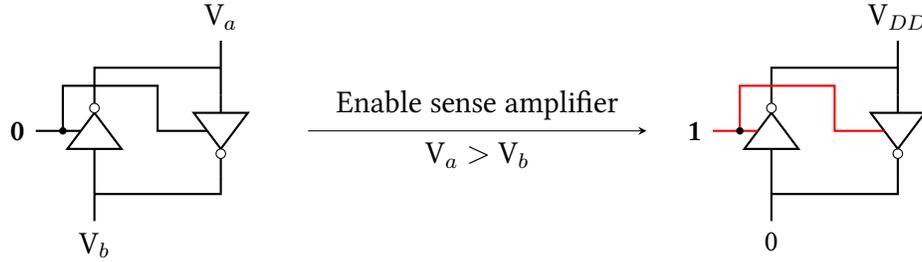


Figure 5: Operation of the sense amplifier

the sense amplifier, and 2) the sequence of states involved in capturing the cell state in the sense amplifier.

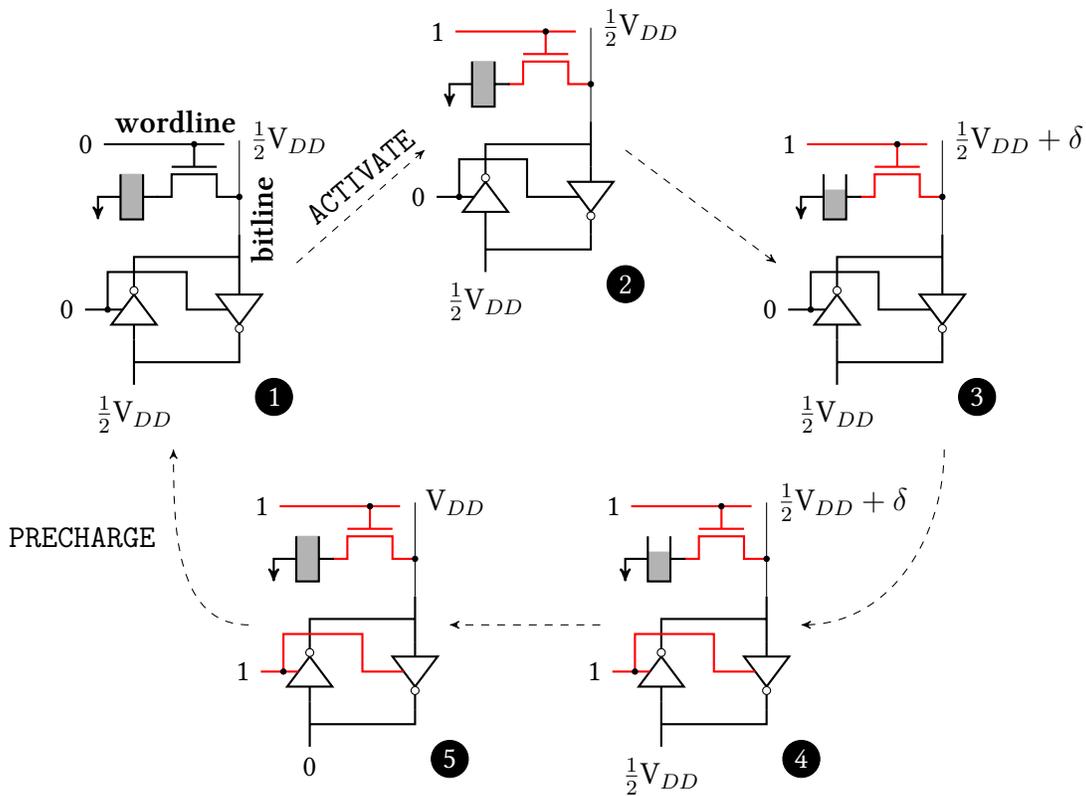


Figure 6: Operation of a DRAM cell and sense amplifier. ① precharged state, ② wordline enable, ③ charge sharing, ④ sense amplifier enable, ⑤ activated state.

As shown in the figure (state ①), the capacitor is connected to an access transistor that acts as a switch between the capacitor and the sense amplifier. The transistor is controller by a wire called *wordline*. The wire that connects the transistor to the top end of the sense amplifier is called *bitline*. In the initial state ①, the wordline is lowered, the sense amplifier is disabled and both ends of the sense amplifier are maintained at a voltage level of $\frac{1}{2}V_{DD}$. We assume that the capacitor is initially fully charged (the operation is similar if the capacitor was initially empty). This state is referred to as the *precharged* state. An access to the cell

is triggered by a command called **ACTIVATE**. Upon receiving an **ACTIVATE**, the corresponding wordline is first raised (state ②). This connects the capacitor to the bitline. In the ensuing phase called *charge sharing* (state ③), charge flows from the capacitor to the bitline, raising the voltage level on the bitline (top end of the sense amplifier) to $\frac{1}{2}V_{DD} + \delta$. After charge sharing, the sense amplifier is enabled (state ④). The sense amplifier detects the difference in voltage levels between its two ends and amplifies the deviation, till it reaches the stable state where the top end is at V_{DD} (state ⑤). Since the capacitor is still connected to the bitline, the charge on the capacitor is also fully restored. We shortly describe how the data can be accessed from the sense amplifier. However, once the access to the cell is complete, the cell is taken back to the original precharged state using the command called **PRECHARGE**. Upon receiving a **PRECHARGE**, the wordline is first lowered, thereby disconnecting the cell from the sense amplifier. Then, the two ends of the sense amplifier are driven to $\frac{1}{2}V_{DD}$ using a precharge unit (not shown in the figure for brevity).

2.2.3. DRAM MAT/Tile: The Open Bitline Architecture

A major goal of DRAM manufacturers is to maximize the density of the DRAM chips while adhering to certain latency constraints (described in Section 2.2.6). There are two costly components in the setup described in the previous section. The first component is the sense amplifier itself. Each sense amplifier is around two orders of magnitude larger than a single DRAM cell [122, 92]. Second, the state of the wordline is a function of the address that is currently being accessed. The logic that is necessary to implement this function (for each cell) is expensive.

In order to reduce the overall cost of these two components, they are shared by many DRAM cells. Specifically, each sense amplifier is shared by a column of DRAM cells. In other words, all the cells in a single column are connected to the same bitline. Similarly, each wordline is shared by a row of DRAM cells. Together, this organization consists of a 2-D array of DRAM cells connected to a row of sense amplifiers and a column of wordline drivers. Figure 7 shows this organization with a 4×4 2-D array.

To further reduce the overall cost of the sense amplifiers and the wordline driver, modern DRAM chips use an architecture called the *open bitline architecture* [31]. This architecture exploits two observations. First, the sense amplifier is wider than the DRAM cells. This difference in width results in a white space near each column of cells. Second, the sense amplifier is symmetric. Therefore, cells can also be connected to the bottom part of the sense amplifier. Putting together these two observations, we can pack twice as many cells in the same area using the open bitline architecture, as shown in Figure 8;

As shown in the figure, a 2-D array of DRAM cells is connected to two rows of sense amplifiers: one on the top and one on the bottom of the array. While all the cells in a given row share a common wordline, half the cells in each row are connected to the top row of sense amplifiers and the remaining half of the cells are connected to the bottom row of sense amplifiers. This tightly packed structure is called a DRAM MAT/Tile [144, 155, 82]. In a modern DRAM chip, each MAT typically is a 512×512 or 1024×1024 array. Multiple MATs are grouped together to form a larger structure called a *DRAM bank*, which we describe next.

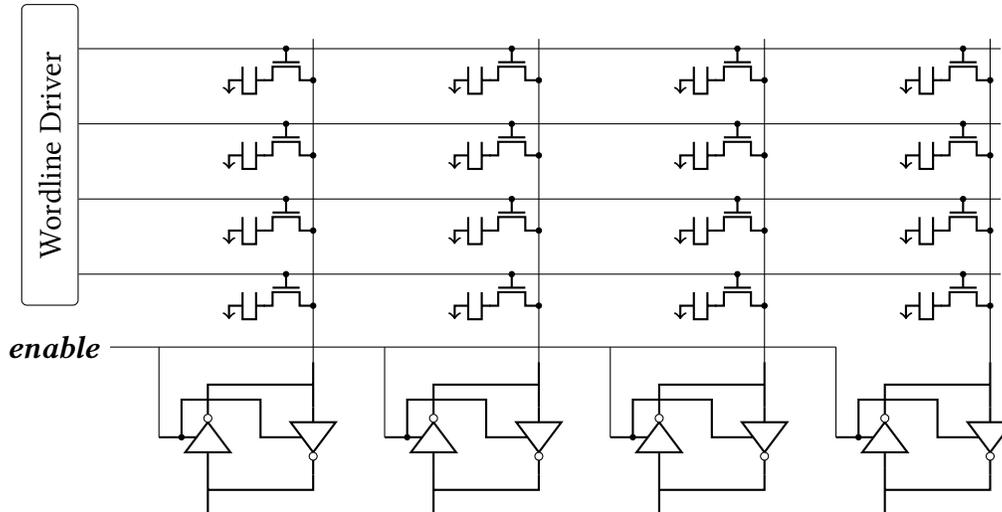


Figure 7: A 2-D array of DRAM cells

2.2.4. DRAM Bank

In most modern commodity DRAM interfaces [65, 66, 83], a DRAM bank is the smallest structure visible to the memory controller. All commands related to data access are directed to a specific bank. Logically, each DRAM bank is a large monolithic structure with a 2-D array of DRAM cells connected to a single set of sense amplifiers (also referred to as a row buffer). For example, in a 2Gb DRAM chip with 8 banks, each bank has 2^{15} rows and each logical row has 8192 DRAM cells. Figure 9 shows this logical view of a bank.

In addition to the MAT, the array of sense amplifiers, and the wordline driver, each bank also consists of some peripheral structures to decode DRAM commands and addresses, and manage the inputs/outputs to the DRAM bank. Specifically, each bank has a *row decoder* to decode the *row address* associated with row-level commands (e.g., **ACTIVATE**). Each data access command (**READ** and **WRITE**) accesses only a part of a DRAM row. Such an individual part is referred to as a *column*. With each data access command, the address of the column to be accessed is provided. This address is decoded by the *column selection logic*. Depending on which column is selected, the corresponding piece of data is communicated between the sense amplifiers and the *bank I/O logic*. The bank I/O logic in turn acts as an interface between the DRAM bank and the *chip-level I/O logic*.

Although the bank can logically be viewed as a single MAT, building a single MAT of a very large size is practically not feasible as it would require very long bitlines and wordlines (leading to very high latencies). Therefore, each bank is physically implemented as a 2-D array of DRAM MATs. Figure 10 shows a physical implementation of the DRAM bank with 4 MATs arranged in a 2×2 array. As shown in the figure, the output of the global row decoder is sent to each row of MATs. The bank I/O logic, also known as the *global sense amplifiers*, are connected to all the MATs through a set of *global bitlines*. As shown in the figure, each vertical collection of MATs has its own column selection logic (CSL) and global bitlines. In a real DRAM chip, the global bitlines run on top of the MATs in a separate metal layer. Data of each column is split equally across a single row of MATs. With this

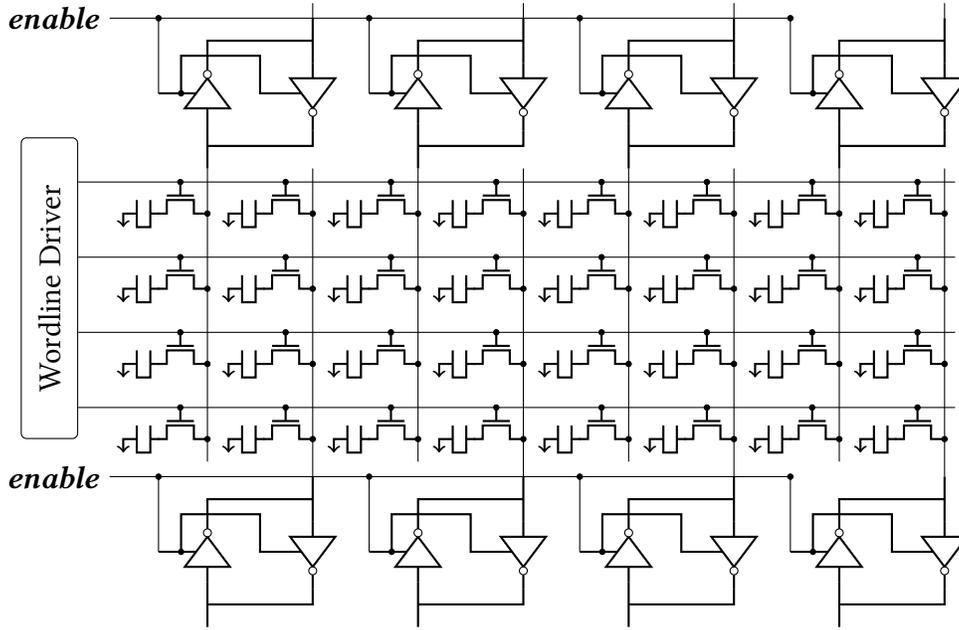


Figure 8: A DRAM MAT/Tile: The open bitline architecture

data organization, each global bitline needs to be connected only to bitlines within one MAT. While prior work [144] has explored routing mechanisms to connect each global bitline to all MATs, such a design incurs high complexity and overhead.

Figure 11 shows the zoomed-in version of a DRAM MAT with the surrounding peripheral logic. Specifically, the figure shows how each column selection logic selects specific sense amplifiers from a MAT and connects them to the global bitlines. We note that the width of the global bitlines for each MAT (typically 8/16) is much smaller than that of the width of the MAT (typically 512/1024). This is because the global bitlines span a much longer distance across the chip and hence have to be wider to ensure signal integrity.

Each DRAM chip consists of multiple banks, as shown in Figure 12. All the banks share the chip’s internal command, address, and data buses. As mentioned before, each bank operates mostly independently (except for operations that involve the shared buses). The chip I/O logic manages the transfer of data to and from the chip’s internal bus to the memory channel. The width of the chip output (typically 8 bits) is much smaller than the output width of each bank (typically 64 bits). Any piece of data accessed from a DRAM bank is first buffered at the chip I/O and sent out on the memory bus 8 bits at a time. With the DDR (double data rate) technology, 8 bits are sent out each half cycle. Therefore, it takes 4 cycles to transfer 64 bits of data from a DRAM chip I/O logic on to the memory channel.

2.2.5. DRAM Commands: Accessing Data from a DRAM Chip

To access a piece of data from a DRAM chip, the memory controller must first identify the location of the data: the bank ID (B), the row address (R) within the bank, and the column address (C) within the row. After identifying these pieces of information, accessing the data involves three steps.

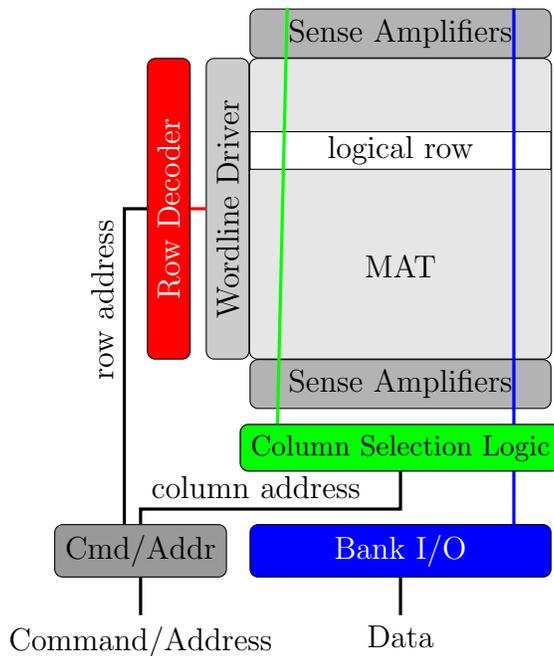


Figure 9: DRAM Bank: Logical view

The first step is to issue a **PRECHARGE** to the bank B . This step prepares the bank for a data access by ensuring that all the sense amplifiers are in the *precharged* state (Figure 6, state ❶). No wordline within the bank is raised in this state.

The second step is to activate the row R that contains the data. This step is triggered by issuing an **ACTIVATE** to bank B with row address R . Upon receiving this command, the corresponding bank feeds its global row decoder with the input R . The global row decoder logic then raises the wordline of the DRAM row corresponding to address R and enables the sense amplifiers connected to that row. This triggers the DRAM cell operation described in Section 2.2.2. At the end of the activate operation, the data from the entire row of DRAM cells is copied to the corresponding array of sense amplifiers.

The third and final step is to access the data from the required column. This is done by issuing a **READ** or **WRITE** command to the bank with the column address C . Upon receiving a **READ** or **WRITE** command, the corresponding address is fed to the column selection logic. The column selection logic then raises the column selection lines (Figure 11) corresponding to address C , thereby connecting those local sense amplifiers to the global sense amplifiers through the global bitlines. For a read access, the global sense amplifiers sense the data from the MAT's local sense amplifiers and transfer that data to the chip's internal bus. For a write access, the global sense amplifiers read the data from the chip's internal bus and force the MAT's local sense amplifiers to the appropriate state.

Not all data accesses require all three steps. Specifically, if the row that is to be accessed is already activated in the corresponding bank, then the first two steps can be skipped and the data can be directly accessed by issuing a **READ** or **WRITE** to the bank. For this reason, the array of sense amplifiers are also referred to as a *row buffer*, and such an access that skips the first two steps is called a *row buffer hit*. Similarly, if the bank is already in the

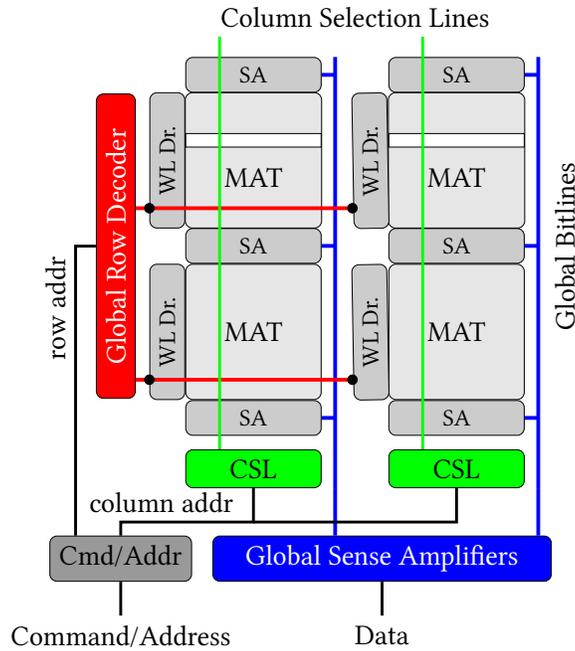


Figure 10: DRAM Bank: Physical implementation. The global bitlines run on top of the MATs in a separate metal layer. (components in the figure are not to scale)

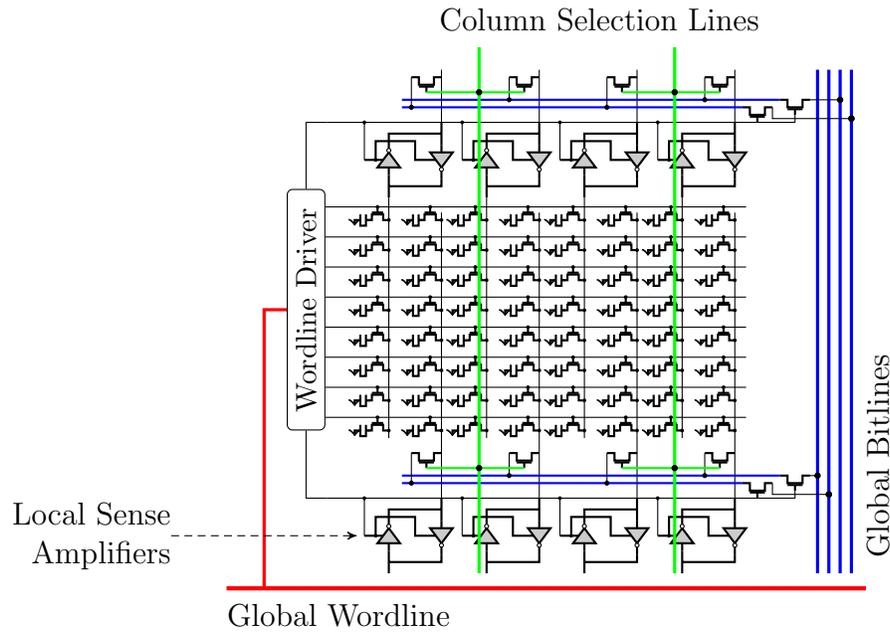


Figure 11: Detailed view of a MAT

precharged state, then the first step can be skipped. Such an access is referred to as a *row buffer miss*. Finally, if a different row is activated within the bank, then all three steps have to be performed. Such an access is referred to as a *row buffer conflict*.

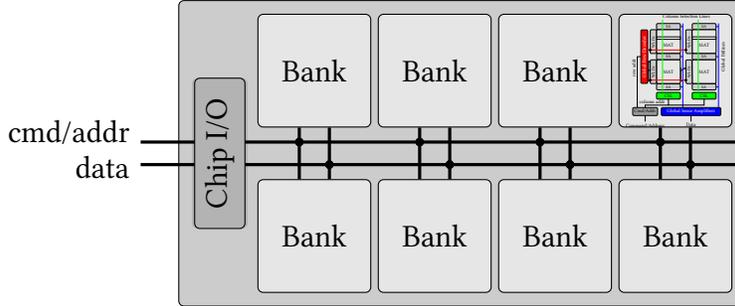


Figure 12: DRAM Chip

2.2.6. DRAM Timing Constraints

Different operations within DRAM consume different amounts of time. Therefore, after issuing a command, the memory controller must wait for a sufficient amount of time before it can issue the next command. Such wait times are managed by pre-specified fixed delays, called the *timing constraints*. Timing constraints essentially dictate the minimum amount of time between two commands issued to the same bank/rank/channel. Table 1 describes some key timing constraints along with their values for the DDR3-1600 interface. The reason as to why these constraints exist is discussed in prior works [82, 92, 32, 62].

Name	Constraint		Description	Value (ns)
tRAS	ACTIVATE	→ PRECHARGE	Time taken to complete a row activation operation in a bank	35
tRCD	ACTIVATE	→ READ/WRITE	Time between an activate command and a column command to a bank	15
tRP	PRECHARGE	→ ACTIVATE	Time taken to complete a precharge operation in a bank	15
tWR	WRITE	→ PRECHARGE	Time taken to ensure that data is safely written to the DRAM cells after a write operation (called <i>write recovery</i>)	15

Table 1: Key DRAM timing constraints, and their values for DDR3-1600 [65]

2.3. DRAM Module

As mentioned before, each READ or WRITE command for a single DRAM chip typically involves only 64 bits. In order to achieve high memory bandwidth, commodity DRAM modules group several DRAM chips (typically 4 or 8) together to form a *rank* of DRAM chips. The idea is to connect all chips of a single rank to the same command and address buses, while providing

each chip with an independent data bus. In effect, all the chips within a rank receive the same command with the same address, making the rank a *logically large* DRAM chip.

Figure 13 shows the logical organization of a DRAM rank. Many commodity DRAM ranks consist of 8 chips with each chip accessing 8 bytes of data in response to each `READ` or `WRITE` command. Therefore, in total, each `READ` or `WRITE` command accesses 64 bytes of data, the typical cache line size in many processors.

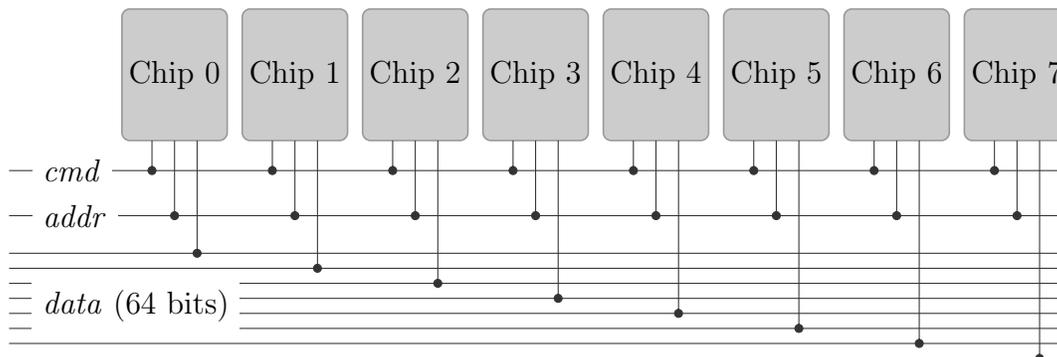


Figure 13: Organization of a DRAM rank

2.4. RowClone: Bulk Copy and Initialization using DRAM

RowClone [133] is a mechanism to perform bulk copy and initialization operations *completely* inside the DRAM. This approach obviates the need to transfer large quantities of data on the memory channel, thereby significantly improving the efficiency of a bulk copy operation. As bulk data initialization (and specifically bulk zeroing) can be viewed as a special case of a bulk copy operation, RowClone can be easily extended to perform such bulk initialization operations with high efficiency.

RowClone consists of two independent mechanisms that exploit several observations about DRAM organization and operation. The first mechanism, called the *Fast Parallel Mode* (FPM), efficiently copies data between two rows of DRAM cells that share the same set of sense amplifiers (i.e., two rows within the same subarray). To copy data from a source row to a destination row within the same subarray, RowClone-FPM first issues an `ACTIVATE` to the source row, immediately followed by an `ACTIVATE` to the destination row. We show that this sequence of back-to-back row activations inside the same subarray results in a data copy from the source row to the destination row. Figure 14 shows the operation of RowClone-FPM on a single cell.

Without any further optimizations, the latency of RowClone-FPM is equivalent to that of two row activations followed by a precharge, which is an order of magnitude faster than existing systems [133].

The second mechanism, called the *Pipelined Serial Mode* (PSM), efficiently copies cache lines between two banks within a module in a pipelined manner. To copy data from a source row in one bank to a destination row in a different bank, RowClone-PSM first activates both the rows. It then uses a newly-proposed command, `TRANSFER`, to copy a single cache

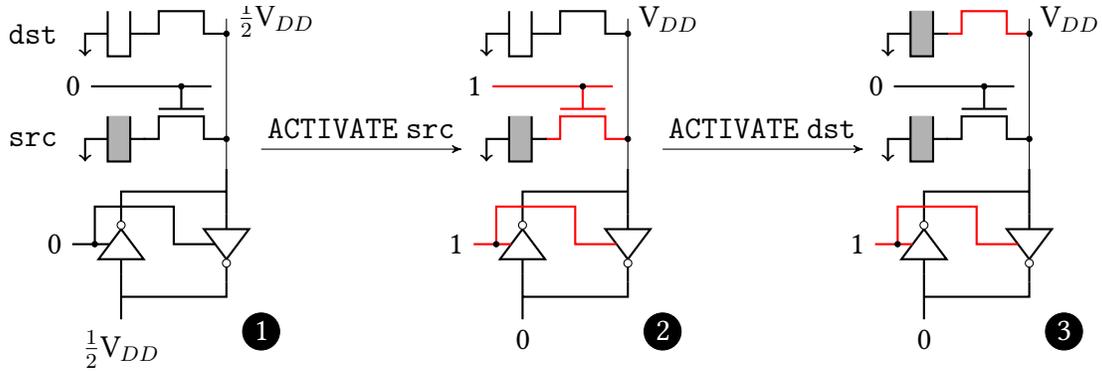


Figure 14: RowClone: Fast Parallel Mode (FPM)

line from the source row *directly* into the destination row, without having to send the data outside the chip. Figure 15 compares the data path of READ, WRITE, and TRANSFER. Although not as fast as FPM, PSM has fewer constraints and hence is more generally applicable [133].

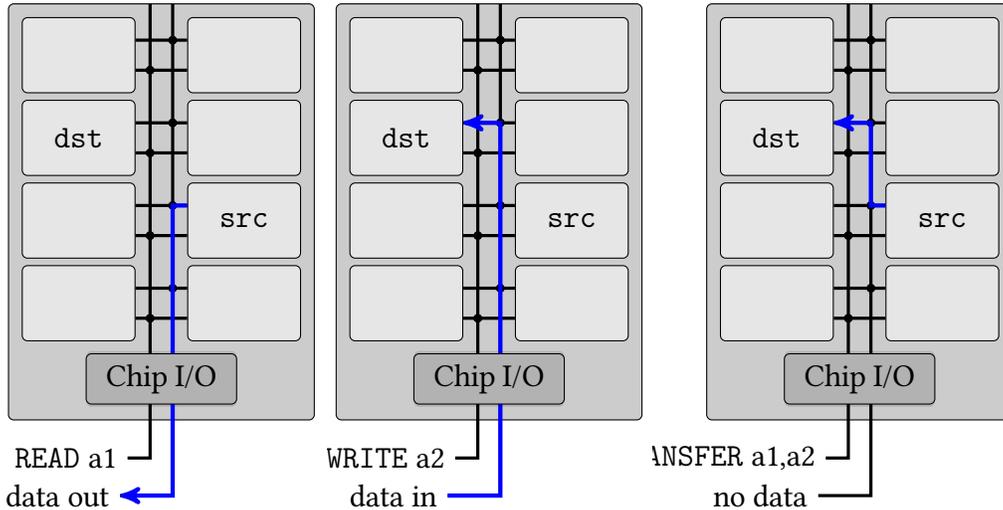


Figure 15: RowClone: Pipelined Serial Mode (PSM)

We refer the reader to [133] for detailed information on RowClone.

3. Ambit: A Bulk Bitwise Execution Engine

In this section, we describe the design and implementation of Ambit, a new mechanism that converts DRAM into a bulk bitwise execution engine with low cost. As mentioned in the introduction, bulk bitwise operations are triggered by many important applications—e.g., bitmap indices [28, 113, 5, 114, 10, 11], databases [97, 98], document filtering [47], DNA processing [21, 150, 18, 77, 111, 129, 17], encryption algorithms [143, 51, 106], graph processing [96], and networking [146]. Accelerating bulk bitwise operations can significantly boost the performance of these applications.

3.1. Ambit-AND-OR

The first component of our mechanism, Ambit-AND-OR, uses the analog nature of the charge sharing phase to perform bulk bitwise AND and OR directly inside the DRAM chip. It specifically exploits two facts about DRAM operation:

1. In a subarray, each sense amplifier is shared by many (typically 512 or 1024) DRAM cells on the same bitline.
2. The final state of the bitline after sense amplification depends primarily on the voltage deviation on the bitline after the charge sharing phase.

Based on these facts, we observe that simultaneously activating three cells, rather than a single cell, results in a *bitwise majority function*—i.e., at least two cells have to be fully charged for the final state to be a logical “1”. We refer to simultaneous activation of three cells (or rows) as *triple-row activation*. We now conceptually describe triple-row activation and how we use it to perform bulk bitwise AND and OR operations.

3.1.1. Triple-Row Activation (TRA)

A triple-row activation (TRA) simultaneously connects a sense amplifier with three DRAM cells on the same bitline. For ease of conceptual understanding, let us assume that the three cells have the same capacitance, the transistors and bitlines behave ideally (no resistance), and the cells start at a fully refreshed state. Then, based on charge sharing principles [69], the bitline deviation at the end of the charge sharing phase of the TRA is:

$$\begin{aligned}\delta &= \frac{k \cdot C_c \cdot V_{DD} + C_b \cdot \frac{1}{2} V_{DD}}{3C_c + C_b} - \frac{1}{2} V_{DD} \\ &= \frac{(2k - 3)C_c}{6C_c + 2C_b} V_{DD}\end{aligned}\tag{1}$$

where, δ is the bitline deviation, C_c is the cell capacitance, C_b is the bitline capacitance, and k is the number of cells in the fully charged state. It is clear that $\delta > 0$ if and only if $2k - 3 > 0$. In other words, the bitline deviation is positive if $k = 2, 3$ and it is negative if $k = 0, 1$. Therefore, we expect the final state of the bitline to be V_{DD} if at least two of the three cells are initially fully charged, and the final state to be 0, if at least two of the three cells are initially fully empty.

Figure 16 shows an example TRA where two of the three cells are initially in the charged state **1**. When the wordlines of all the three cells are raised simultaneously **2**, charge sharing results in a positive deviation on the bitline. Therefore, after sense amplification **3**, the sense amplifier drives the bitline to V_{DD} , and as a result, fully charges all the three cells.¹

If A , B , and C represent the logical values of the three cells, then the final state of the bitline is $AB + BC + CA$ (the *bitwise majority function*). Importantly, we can rewrite this expression as $C(A + B) + \overline{C}(AB)$. In other words, by controlling the value of the cell C , we can use TRA to execute a *bitwise AND* or *bitwise OR* of the cells A and B . Since activation

¹Modern DRAMs use an open-bitline architecture [90, 31, 100, 69] (Section 2.2.3), where cells are also connected to bitline. The three cells in our example are connected to the bitline. However, based on the duality principle of Boolean algebra [138], i.e., not (A and B) \equiv (not A) or (not B), TRA works seamlessly even if all the three cells are connected to bitline.

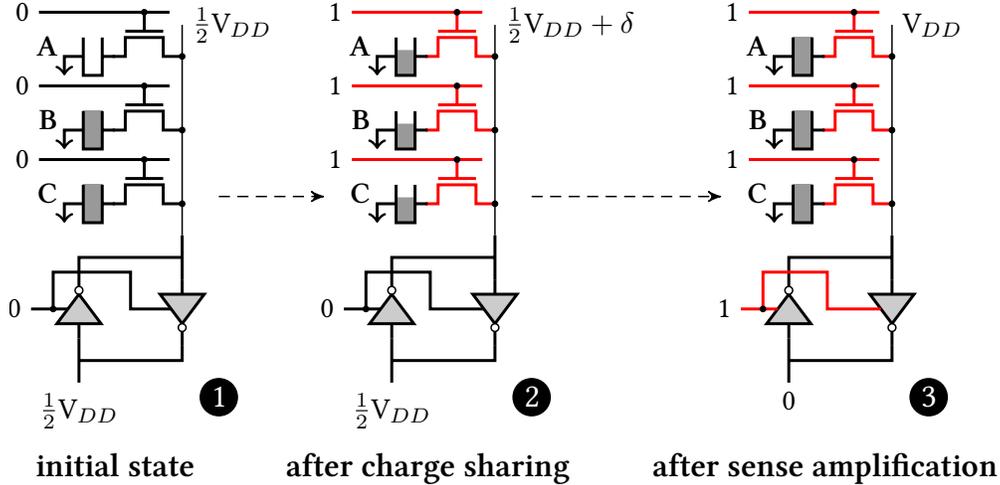


Figure 16: Triple-row activation (TRA)

is a row-level operation in DRAM, TRA operates on an *entire row* of DRAM cells and sense amplifiers, thereby enabling a multi-kilobyte-wide bitwise AND/OR of two rows.

3.1.2. Making TRA Work

There are five potential issues with TRA that we need to resolve for it to be implementable in a real DRAM design.

1. When simultaneously activating three cells, the deviation on the bitline may be smaller than when activating only one cell. This may lengthen sense amplification or worse, the sense amplifier may detect the wrong value.
2. Equation 1 assumes that all cells have the same capacitance, and that the transistors and bitlines behave ideally. However, due to process variation, these assumptions are not true in real designs [91, 90, 32]. This can affect the reliability of TRA, and thus the correctness of its results.
3. As shown in Figure 16 (state 3), TRA overwrites the data of all the three cells with the final result value. In other words, TRA overwrites all source cells, thereby destroying their original values.
4. Equation 1 assumes that the cells involved in a TRA are either fully-charged or fully-empty. However, DRAM cells leak charge over time [101]. If the cells involved have leaked significantly, TRA may not operate as expected.
5. Simultaneously activating three *arbitrary* rows inside a DRAM subarray requires the memory controller and the row decoder to simultaneously communicate and decode three row addresses. This introduces a large cost on the address bus and the row decoder, potentially tripling these structures, if implemented naïvely.

We address the first two issues by performing rigorous circuit simulations of TRA. Our results confirm that TRA works as expected (Section 6). In Sections 3.1.3 and 3.1.4, we propose a simple implementation of Ambit-AND-OR that addresses *all* of the last three issues at low cost.

3.1.3. Implementation of Ambit-AND-OR

To solve issues 3, 4, and 5 described in Section 3.1.2, our implementation of Ambit reserves a set of *designated rows* in each subarray that are used to perform TRAs. These designated rows are chosen statically at *design time*. To perform a bulk bitwise AND or OR operation on two arbitrary source rows, our mechanism *first copies* the data of the source rows into the designated rows and performs the required TRA on the designated rows. As an example, to perform a bitwise AND/OR of two rows A and B, and store the result in row R, our mechanism performs the following steps.

1. *Copy* data of row A to designated row T0
2. *Copy* data of row B to designated row T1
3. *Initialize* designated row T2 to 0
4. *Activate* designated rows T0, T1, and T2 simultaneously
5. *Copy* data of row T0 to row R

Let us understand how this implementation addresses the last three issues described in Section 3.1.2. First, by performing the TRA on the designated rows, and *not* directly on the source data, our mechanism avoids overwriting the source data (issue 3). Second, each copy operation refreshes the cells of the destination row by accessing the row [101]. Also, each copy operation takes five-six orders of magnitude lower latency (100 ns—1 μ s) than the refresh interval (64 ms). Since these copy operations (Steps 1 and 2 above) are performed *just before* the TRA, the rows involved in the TRA are very close to the fully-refreshed state just before the TRA operation (issue 4). Finally, since the *designated* rows are chosen statically at design time, the Ambit controller uses a reserved address to communicate the TRA of a *pre-defined* set of three designated rows. To this end, Ambit reserves a set of row addresses *just* to trigger TRAs. For instance, in our implementation to perform a TRA of designated rows T0, T1, and T2 (Step 4, above), the Ambit controller simply issues an **ACTIVATE** with the reserved address B12 (see Section 4.1 for a full list of reserved addresses). The row decoder maps B12 to *all* the three wordlines of the designated rows T0, T1, and T2. This mechanism requires *no* changes to the address bus and significantly reduces the cost and complexity of the row decoder compared to performing TRA on three *arbitrary* rows (issue 5).

3.1.4. Fast Row Copy and Initialization Using RowClone

Our mechanism needs three row copy operations and one row initialization operation. These operations, if performed naively, can nullify the benefits of Ambit, as a row copy or row initialization performed using the memory controller incurs high latency [133, 31]. Fortunately, a recent work, RowClone [133], proposes two techniques to efficiently copy data between rows *directly within* DRAM. The first technique, RowClone-FPM (Fast Parallel Mode), copies data within a subarray by issuing two back-to-back **ACTIVATE**s to the source row and the destination row. This operation takes only 80 ns [133]. The second technique, RowClone-PSM (Pipelined Serial Mode), copies data between two banks by using the internal DRAM bus. Although RowClone-PSM is faster and more efficient than copying data using the memory controller, it is significantly slower than RowClone-FPM.

Ambit relies on using RowClone-FPM for most of the copy operations. To enable this, we propose three ideas. First, to allow Ambit to perform the initialization operation using RowClone-FPM, we reserve two *control* rows in each subarray, **C0** and **C1**. **C0** is initialized to all zeros and **C1** is initialized to all ones. Depending on the operation to be performed, bitwise AND or OR, Ambit copies the data from **C0** or **C1** to the appropriate designated row using RowClone-FPM. Second, we reserve separate designated rows in *each* subarray. This allows each subarray to perform bulk bitwise AND/OR operations on the rows that belong to that subarray by using RowClone-FPM for all the required copy operations. Third, to ensure that bulk bitwise operations are predominantly performed between rows inside the *same* subarray, we rely on 1) an accelerator API that allows applications to specify bitvectors that are likely to be involved in bitwise operations, and 2) a driver that maps such bitvectors to the *same* subarray (described in Section 5.2). With these changes, Ambit can use RowClone-FPM for a significant majority of the bulk copy operations, thereby ensuring high performance for the bulk bitwise operations.

A recent work, Low-cost Interlinked Subarrays (LISA) [31], proposes a mechanism to efficiently copy data across subarrays in the same bank. LISA uses a row of isolation transistors next to the sense amplifier to control data transfer across two subarrays. LISA can potentially benefit Ambit by improving the performance of bulk copy operations. However, as we will describe in Section 3.2, Ambit-NOT also adds transistors near the sense amplifier, posing some challenges in integrating LISA and Ambit. Therefore, we leave the exploration of using LISA to speed up Ambit as part of future work.

3.2. Ambit-NOT

Ambit-NOT exploits the fact that at the end of the sense amplification process, the voltage level of the $\overline{\text{bitline}}$ represents the negated logical value of the cell. Our key idea to perform bulk bitwise NOT in DRAM is to transfer the data on the $\overline{\text{bitline}}$ to a cell that can *also* be connected to the bitline. For this purpose, we introduce the *dual-contact cell* (shown in Figure 17). A dual-contact cell (DCC) is a DRAM cell with two transistors (a 2T-1C cell similar to the one described in [67, 103]). Figure 17 shows a DCC connected to a sense amplifier. In a DCC, one transistor connects the cell capacitor to the bitline and the other transistor connects the cell capacitor to the $\overline{\text{bitline}}$. We refer to the wordline that controls the capacitor-bitline connection as the *d-wordline* (or data wordline) and the wordline that controls the capacitor- $\overline{\text{bitline}}$ connection as the *n-wordline* (or negation wordline). The layout of the dual-contact cell is similar to Lu et al.’s migration cell [103].

Figure 18 shows the steps involved in transferring the negated value of a *source* cell on to the DCC connected to the same bitline (i.e., sense amplifier) ❶. Our mechanism first activates the source cell ❷. The activation drives the bitline to the data value corresponding to the source cell, V_{DD} in this case and the $\overline{\text{bitline}}$ to the negated value, i.e., 0 ❸. In this *activated* state, our mechanism activates the *n-wordline*. Doing so enables the transistor that connects the DCC to the $\overline{\text{bitline}}$ ❹. Since the $\overline{\text{bitline}}$ is already at a stable voltage level of 0, it overwrites the value in the DCC capacitor with 0, thereby copying the negated value of the source cell into the DCC. After this, our mechanism *precharges* the bank, and then copies the negated value from the DCC to the destination cell using RowClone.

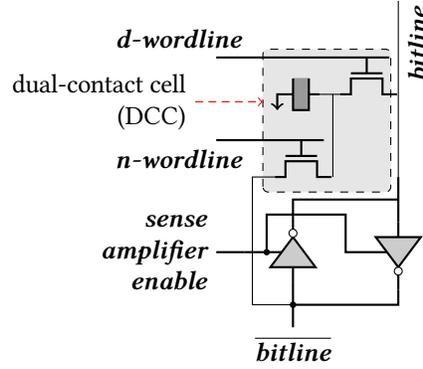


Figure 17: A dual-contact cell connected to a sense amplifier

Implementation of Ambit-NOT. Based on Lu et al.’s [103] layout, the cost of each row of DCC is the same as two regular DRAM rows. Similar to the designated rows used for Ambit-AND-OR (Section 3.1.3), the Ambit controller uses reserved row addresses to control the *d-wordlines* and *n-wordlines* of the DCC rows—e.g., in our implementation, address B5 maps to the *n-wordline* of the DCC row (Section 4.1). To perform a bitwise NOT of row A and store the result in row R, the Ambit controller performs the following steps.

1. *Activate* row A
2. *Activate* *n-wordline* of DCC (address B5)
3. *Precharge* the bank.
4. *Copy* data from *d-wordline* of DCC to row R (RowClone)

4. Ambit: Full Design and Implementation

In this section, we describe our implementation of Ambit by integrating Ambit-AND-OR and Ambit-NOT. First, both Ambit-AND-OR and Ambit-NOT reserve a set of rows in each subarray and a set of addresses that map to these rows. We present the full set of reserved addresses and their mapping in detail (Section 4.1). Second, we introduce a new primitive called AAP (ACTIVATE-ACTIVATE-PRECHARGE) that the Ambit controller uses to execute various bulk bitwise operations (Section 4.2). Third, we describe an optimization that lowers the latency of the AAP primitive, further improving the performance of Ambit

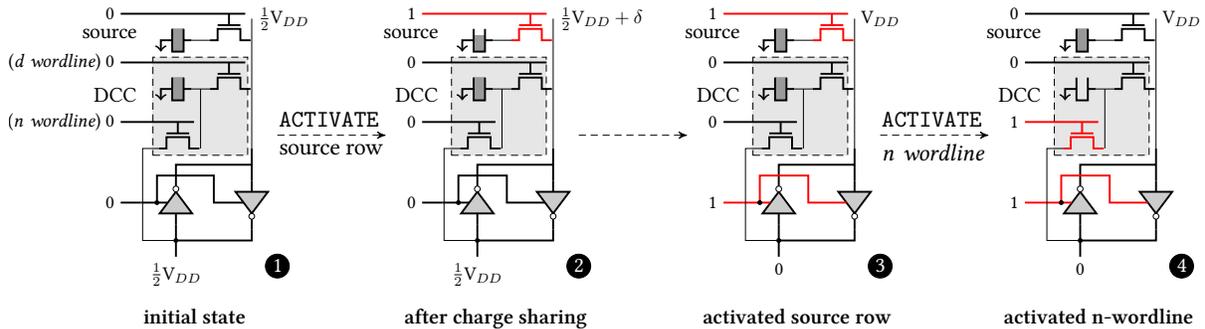


Figure 18: Bitwise NOT using a dual-contact cell

(Section 4.3). Fourth, we describe how we integrate Ambit with the system stack (Section 5). Finally, we evaluate the hardware cost of Ambit (Section 5.6).

4.1. Row Address Grouping

Our implementation divides the space of row addresses in each subarray into three distinct groups (Figure 19): 1) **Bitwise group**, 2) **Control group**, and 3) **Data group**.

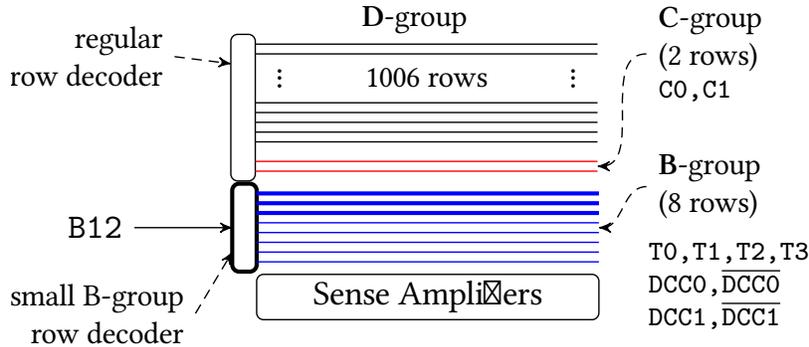


Figure 19: Row address grouping in a subarray. The figure shows how the B-group row decoder (Section 4.3) simultaneously activates rows T_0 , T_1 , and T_2 with a single address B_{12} .

The *B-group* (or the *bitwise group*) corresponds to the designated rows used to perform bulk bitwise AND/OR operations (Section 3.1.3) and the dual-contact rows used to perform bulk bitwise NOT operations (Section 3.2). Minimally, Ambit requires three designated rows (to perform triple row activations) and one row of dual-contact cells in each subarray. However, to reduce the number of copy operations required by certain bitwise operations (like `xor` and `xnor`), we design each subarray with four designated rows, namely T_0 — T_3 , and two rows of dual-contact cells, one on each side of the row of sense amplifiers.² We refer to the *d-wordlines* of the two DCC rows as DCC_0 and DCC_1 , and the corresponding *n-wordlines* as $\overline{DCC_0}$ and $\overline{DCC_1}$. The B-group contains 16 reserved addresses: B_0 — B_{15} . Table 2 lists the mapping between the 16 addresses and the wordlines. The first eight addresses individually activate each of the 8 wordlines in the group. Addresses B_{12} — B_{15} activate three wordlines simultaneously. Ambit uses these addresses to trigger triple-row activations. Finally, addresses B_8 — B_{11} activate two wordlines. Ambit uses these addresses to copy the result of an operation simultaneously to two rows. This is useful for `xor/xnor` operations to *simultaneously* negate a row of source data and also copy the source row to a designated row. Note that this is just an example implementation of Ambit and a real implementation may use more designated rows in the B-group, thereby enabling more complex bulk bitwise operations with fewer copy operations.

²Each `xor/xnor` operation involves multiple `and`, `or`, and `not` operations. We use the additional designated row and the DCC row to store *intermediate* results computed as part of the `xor/xnor` operation (see Figure 20c).

Addr.	Wordline(s)	Addr.	Wordline(s)
B0	T0	B8	$\overline{\text{DCC0}}$, T0
B1	T1	B9	$\overline{\text{DCC1}}$, T1
B2	T2	B10	T2, T3
B3	T3	B11	T0, T3
B4	DCC0	B12	T0, T1, T2
B5	$\overline{\text{DCC0}}$	B13	T1, T2, T3
B6	DCC1	B14	DCC0, T1, T2
B7	$\overline{\text{DCC1}}$	B15	DCC1, T0, T3

Table 2: Mapping of B-group addresses to corresponding activated wordlines

The *C-group* (or the *control* group) contains the two pre-initialized rows for controlling the bitwise AND/OR operations (Section 3.1.4). Specifically, this group contains two addresses: C0 (row with all zeros) and C1 (row with all ones).

The *D-group* (or the *data* group) corresponds to the rows that store regular data. This group contains all the addresses that are neither in the *B-group* nor in the *C-group*. Specifically, if each subarray contains 1024 rows, then the *D-group* contains 1006 addresses, labeled D0—D1005. Ambit exposes only the D-group addresses to the software stack. To ensure that the software stack has a contiguous view of memory, the Ambit controller interleaves the row addresses such that the D-group addresses across all subarrays are mapped contiguously to the processor’s physical address space.

With these groups, the Ambit controller can use the *existing* DRAM interface to communicate all variants of **ACTIVATE** to the Ambit chip *without* requiring new commands. Depending on the address group, the Ambit DRAM chip internally processes each **ACTIVATE** appropriately. For instance, by just issuing an **ACTIVATE** to address B12, the Ambit controller triggers a triple-row activation of T0, T1, and T2. We now describe how the Ambit controller uses this row address mapping to perform bulk bitwise operations.

4.2. Executing Bitwise Ops: The AAP Primitive

Let us consider the operation, $D_k = \text{not } D_i$. To perform this bitwise-NOT operation, the Ambit controller sends the following sequence of commands.

1. **ACTIVATE** D_i ; 2. **ACTIVATE** B5; 3. **PRECHARGE**;
4. **ACTIVATE** B4; 5. **ACTIVATE** D_k ; 6. **PRECHARGE**;

The first three steps are the same as those described in Section 3.2. These three operations copy the negated value of row D_i into the DCC0 row (as described in Figure 18). Step 4 activates DCC0, the *d-wordline* of the first DCC row, transferring the negated source data onto the bitlines. Step 5 activates the destination row, copying the data on the bitlines, i.e., the negated source data, to the destination row. Step 6 prepares the array for the next access by issuing a **PRECHARGE**.

The bitwise-NOT operation consists of two steps of **ACTIVATE-ACTIVATE-PRECHARGE** operations. We refer to this sequence as the **AAP** primitive. Each **AAP** takes two addresses as

input. AAP (addr1, addr2) corresponds to the following sequence of commands:

```
ACTIVATE addr1; ACTIVATE addr2; PRECHARGE;
```

Logically, an AAP operation copies the result of the row activation of the first address (addr1) to the row(s) mapped to the second address (addr2).

Most bulk bitwise operations mainly involve a sequence of AAP operations. In a few cases, they require a regular ACTIVATE followed by a PRECHARGE, which we refer to as AP. AP takes one address as input. AP (addr) maps to the following two commands:

```
ACTIVATE addr; PRECHARGE;
```

Figure 20 shows the steps taken by the Ambit controller to execute three bulk bitwise operations: `and`, `nand`, and `xor`.

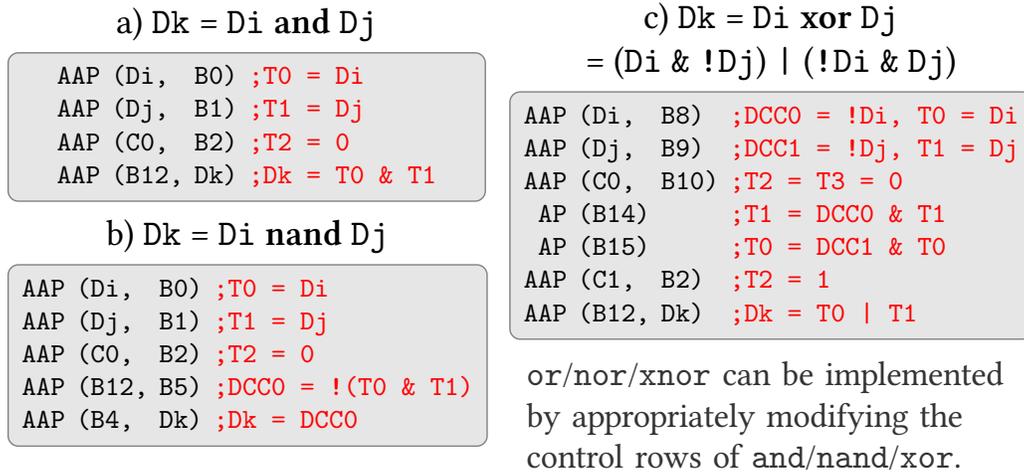


Figure 20: Command sequences for different bitwise operations

Let us consider the `and` operation, $D_k = D_i \text{ and } D_j$, shown in Figure 20a. The four AAP operations directly map to the steps described in Section 3.1.3. The first AAP copies the first source row (D_i) into the designated row T_0 . Similarly, the second AAP copies the second source row D_j to row T_1 , and the third AAP copies the control row “0” to row T_2 (to perform a bulk bitwise AND). Finally, the last AAP 1) issues an ACTIVATE to address B12, which simultaneously activates the rows T_0 , T_1 , and T_2 , resulting in an `and` operation of the rows T_0 and T_1 , 2) issues an ACTIVATE to D_k , which copies the result of the `and` operation to the destination row D_k , and 3) precharges the bank to prepare it for the next access.

While each bulk bitwise operation involves multiple copy operations, this copy overhead can be reduced by applying standard compilation techniques. For instance, accumulation-like operations generate intermediate results that are immediately consumed. An optimization like dead-store elimination may prevent these values from being copied unnecessarily. Our evaluations (Section 8) take into account the overhead of the copy operations *without* such optimizations.

4.3. Accelerating AAP with a Split Row Decoder

The latency of executing any bulk bitwise operation using Ambit depends on the latency of the AAP primitive. The latency of the AAP in turn depends on the latency of `ACTIVATE`, i.e., t_{RAS} , and the latency of `PRECHARGE`, i.e., t_{RP} . The naïve approach to execute an AAP is to perform the three operations serially. Using this approach, the latency of AAP is $2t_{\text{RAS}} + t_{\text{RP}}$ (80 ns for DDR3-1600 [64]). While even this naïve approach offers better throughput and energy efficiency than existing systems (not shown here), we propose a simple optimization that significantly reduces the latency of AAP.

Our optimization is based on two observations. First, the second `ACTIVATE` of an AAP is issued to an already *activated bank*. As a result, this `ACTIVATE` does *not* require *full* sense amplification, which is the dominant portion of t_{RAS} [90, 91, 55]. This enables the opportunity to reduce the latency for the second `ACTIVATE` of each AAP. Second, when we examine all the bitwise operations in Figure 20, with the exception of one AAP in `nand`, we find that *exactly one* of the two `ACTIVATE`s in each AAP is to a *B-group* address. This enables the opportunity to use a *separate* decoder for *B-group* addresses, thereby overlapping the latency of the two row activations in each AAP.

To exploit both of these observations, our mechanism splits the row decoder into two parts. The first part decodes all *C/D-group* addresses and the second smaller part decodes *only B-group* addresses. Such a split allows the subarray to *simultaneously* decode a *C/D-group* address along with a *B-group* address. When executing an AAP, the Ambit controller issues the second `ACTIVATE` of an AAP after the first activation has sufficiently progressed. This forces the sense amplifier to overwrite the data of the second row to the result of the first activation. This mechanism allows the Ambit controller to significantly overlap the latency of the two `ACTIVATE`s. This approach is similar to the inter-segment copy operation used by Tiered-Latency DRAM [92]. Based on SPICE simulations, our estimate of the latency of executing the back-to-back `ACTIVATE`s is only 4 ns larger than t_{RAS} . For DDR3-1600 (8-8-8) timing parameters [64], this optimization reduces the latency of AAP from 80 ns to 49 ns.

Since only addresses in the *B-group* are involved in triple-row activations, the complexity of simultaneously raising three wordlines is restricted to the small *B-group* decoder. As a result, the split row decoder also reduces the complexity of the changes Ambit introduces to the row decoding logic.

5. Integrating Ambit with the System

Ambit can be plugged in as an I/O (e.g., PCIe) device and interfaced with the CPU using a device model similar to other accelerators (e.g., GPU). While this approach is simple, as described in previous sections, the address and command interface of Ambit is *exactly the same* as that of commodity DRAM. This enables the opportunity to directly plug Ambit onto the system memory bus and control it using the memory controller. This approach has several benefits. First, applications can *directly* trigger Ambit operations using CPU instructions rather than going through a device API, which incurs additional overhead. Second, since the CPU can *directly* access Ambit memory, there is no need to copy data between the CPU memory and the accelerator memory. Third, existing cache coherence

protocols can be used to keep Ambit memory and the on-chip cache coherent. To plug Ambit onto the system memory bus, we need additional support from the rest of the system stack, which we describe in this section.

5.1. ISA Support

To enable applications to communicate occurrences of bulk bitwise operations to the processor, we introduce new instructions of the form,

`bbop dst, src1, [src2], size`

where `bbop` is the bulk bitwise operation, `dst` is the destination address, `src1` and `src2` are the source addresses, and `size` denotes the length of operation in bytes. Note that `size` must be a multiple of DRAM row size. For bitvectors that are not a multiple of DRAM row size, we assume that the application will appropriately pad them with dummy data, or perform the residual (sub-row-sized) operations using the CPU.

5.2. Ambit API/Driver Support

For Ambit to provide significant performance benefit over existing systems, it is critical to ensure that most of the required copy operations are performed using RowClone-FPM, i.e., the source rows and the destination rows involved in bulk bitwise operations are present in the same DRAM subarray. To this end, we expect the manufacturer of Ambit to provide 1) an API that enables applications to specify bitvectors that are likely to be involved in bitwise operations, and 2) a driver that is aware of the internal mapping of DRAM rows to subarrays and maps the bitvectors involved in bulk bitwise operations to the same DRAM subarray. Note that for a large bitvector, Ambit does *not* require the entire bitvector to fit inside a *single* subarray. Rather, each bitvector can be interleaved across multiple subarrays such that the corresponding portions of each bitvector are in the same subarray. Since each subarray contains over 1000 rows to store application data, an application can map hundreds of *large* bitvectors to Ambit, such that the copy operations required by *all* the bitwise operations across all these bitvectors can be performed efficiently using RowClone-FPM.

5.3. Implementing the `bbop` Instructions

Since all Ambit operations are row-wide, Ambit requires the source and destination rows to be row-aligned and the size of the operation to be a multiple of the size of a DRAM row. The microarchitecture implementation of a `bbop` instruction checks if each instance of the instruction satisfies this constraint. If so, the CPU sends the operation to the memory controller, which completes the operation using Ambit. Otherwise, the CPU executes the operation itself.

5.4. Maintaining On-chip Cache Coherence

Since both CPU and Ambit can access/modify data in memory, before performing any Ambit operation, the memory controller must 1) flush any dirty cache lines from the source rows, and 2) invalidate any cache lines from destination rows. Such a mechanism is already required by Direct Memory Access (DMA) [34], which is supported by most modern processors, and also by recently proposed mechanisms [133, 59, 25, 27]. As Ambit operations are always row-wide, we can use structures like the Dirty-Block Index [131] to speed up flushing dirty data. Our mechanism invalidates the cache lines of the destination rows in parallel with the Ambit operation. Other recently-proposed techniques like LazyPIM [25] and CoNDA [26] can also be employed with Ambit.

5.5. Error Correction and Data Scrambling

In DRAM modules that support Error Correction Code (ECC), the memory controller must first read the data and ECC to verify data integrity. Since Ambit reads and modifies data directly in memory, it does not work with the existing ECC schemes (e.g., SECDED [50]). To support ECC with Ambit, we need an ECC scheme that is homomorphic [126] over all bitwise operations, i.e., $ECC(A \text{ and } B) = ECC(A) \text{ and } ECC(B)$, and similarly for other bitwise operations. The only scheme that we are aware of that has this property is triple modular redundancy (TMR) [105], wherein $ECC(A) = AA$. The design of lower-overhead ECC schemes that are homomorphic over all bitwise operations is an open problem. For the same reason, Ambit does *not* work with data scrambling mechanisms that pseudo-randomly modify the data written to DRAM [40]. Note that these challenges are also present in any mechanism that *interprets* data directly in memory (e.g., the Automata Processor [37, 140]). We leave the evaluation of Ambit with TMR and exploration of other ECC and data scrambling schemes to future work.

5.6. Ambit Hardware Cost

As Ambit largely exploits the structure and operation of existing DRAM design, we estimate its hardware cost in terms of the overhead it imposes on top of today’s DRAM chip and memory controller.

5.6.1. Ambit Chip Cost

In addition to support for RowClone, Ambit has only two changes on top of the existing DRAM chip design. First, it requires the row decoding logic to distinguish between the *B-group* addresses and the remaining addresses. Within the *B-group*, it must implement the mapping described in Table 2. As the *B-group* contains only 16 addresses, the complexity of the changes to the row decoding logic is low. The second source of cost is the implementation of the dual-contact cells (DCCs). In our design, each sense amplifier has only one DCC on each side, and each DCC has two wordlines associated with it. In terms of area, each DCC row costs roughly two DRAM rows, based on estimates from Lu et al. [103]. We estimate

the overall storage cost of Ambit to be roughly 8 DRAM rows per subarray—for the four designated rows and the DCC rows (< 1% of DRAM chip area).

5.6.2. Ambit Controller Cost

On top of the existing memory controller, the Ambit controller must statically store 1) information about different address groups, 2) the timing of different variants of the `ACTIVATE`, and 3) the sequence of commands required to complete different bitwise operations. When Ambit is plugged onto the system memory bus, the controller can interleave the various AAP operations in the bitwise operations with other regular memory requests from different applications. For this purpose, the Ambit controller must also track the status of on-going bitwise operations. We expect the overhead of these additional pieces of information to be small compared to the benefits enabled by Ambit.

5.6.3. Ambit Testing Cost

Testing Ambit chips is similar to testing regular DRAM chips. In addition to the regular DRAM rows, the manufacturer must test if the TRA operations and the DCC rows work as expected. In each subarray with 1024 rows, these operations concern only 8 DRAM rows and 16 addresses of the B-group. In addition, all these operations are triggered using the `ACTIVATE` command. Therefore, we expect the overhead of testing an Ambit chip on top of testing a regular DRAM chip to be low.

When a component is found to be faulty during testing, DRAM manufacturers use a number of techniques to improve the overall yield; The most prominent among them is using spare rows to replace faulty DRAM rows. Similar to some prior works [133, 92, 91, 82], Ambit requires faulty rows to be mapped to spare rows within the *same* subarray. Note that, since Ambit preserves the existing DRAM command interface, an Ambit chip that fails during testing can still be shipped as a regular DRAM chip. This significantly reduces any potential negative impact of Ambit-specific failures on overall DRAM yield.

6. Circuit-level SPICE Simulations

We use SPICE simulations to confirm that Ambit works reliably. Of the two components of Ambit, our SPICE results show that Ambit-NOT *always* works as expected and is *not* affected by process variation. This is because, Ambit-NOT operation is very similar to existing DRAM operation (Section 3.2). On the other hand, Ambit-AND-OR requires triple-row activation, which involves charge sharing between three cells on a bitline. As a result, it can be affected by process variation in various circuit components.

To study the effect of process variation on TRA, our SPICE simulations model variation in *all* the components in the subarray (cell capacitance; transistor length, width, resistance; bitline/wordline capacitance and resistance; voltage levels). We implement the sense amplifier using 55nm DDR3 model parameters [13], and PTM low-power transistor models [9, 156].

We use cell/transistor parameters from the Rambus power model [13] (cell capacitance = 22fF; transistor width/height = 55nm/85nm).³

We first identify the worst case for TRA, wherein every component has process variation that works toward making TRA fail. Our results show that even in this extremely adversarial scenario, TRA works reliably for up to $\pm 6\%$ variation in each component.

In practice, variations across components are not so highly correlated. Therefore, we use Monte-Carlo simulations to understand the practical impact of process variation on TRA. We increase the amount of process variation from $\pm 5\%$ to $\pm 25\%$ and run 100,000 simulations for each level of process variation. Table 3 shows the percentage of iterations in which TRA operates incorrectly for each level of variation.

Variation	$\pm 0\%$	$\pm 5\%$	$\pm 10\%$	$\pm 15\%$	$\pm 20\%$	$\pm 25\%$
% Failures	0.00%	0.00%	0.29%	6.01%	16.36%	26.19%

Table 3: Effect of process variation on TRA

Two conclusions are in order. First, as expected, up to $\pm 5\%$ variation, there are zero errors in TRA. Second, even with $\pm 10\%$ and $\pm 15\%$ variation, the percentage of erroneous TRAs across 100,000 iterations each is just 0.29% and 6.01%. These results show that Ambit is reliable even in the presence of significant process variation.

The effect of process variation is expected to get worse with smaller technology nodes [68]. However, as Ambit largely uses the existing DRAM structure and operation, many techniques used to combat process variation in existing chips can be used for Ambit as well (e.g., spare rows or columns). In addition, as described in Section 5.6.3, Ambit chips that fail testing only for TRA can potentially be shipped as regular DRAM chips, thereby alleviating the impact of TRA failures on overall DRAM yield, and thus cost.

Note that if an application can tolerate errors in computation, as observed by prior works [107, 104, 153, 142], then Ambit can be used as an approximate in-DRAM computation substrate even in the presence of significant process variation. Future work can potentially explore such an approximate Ambit substrate.

7. Analysis of Ambit’s Throughput & Energy

We compare the raw throughput of bulk bitwise operations using Ambit to a multi-core Intel Skylake CPU [1], an NVIDIA GeForce GTX 745 GPU [6], and processing in the logic layer of an HMC 2.0 [8] device. The Intel CPU has 4 cores with Advanced Vector eXtensions [61], and two 64-bit DDR3-2133 channels. The GTX 745 contains 3 streaming multi-processors, each with 128 CUDA cores [99], and one 128-bit DDR3-1800 channel. The HMC 2.0 device consists of 32 vaults each with 10 GB/s bandwidth. We use two Ambit configurations: *Ambit* that integrates our mechanism into a regular DRAM module with 8 banks, and

³In DRAM, temperature affects mainly cell leakage [91, 100, 152, 125, 101, 117, 32, 56, 80, 78, 30, 44, 76]. As TRA is performed on cells that are almost fully-refreshed, we do not expect temperature to affect TRA.

Ambit-3D that extends a 3D-stacked DRAM similar to HMC with support for Ambit. For each bitwise operation, we run a microbenchmark that performs the operation repeatedly for many iterations on large input vectors (32 MB), and measure the throughput of the operation. Figure 21 plots the results of this experiment for the five systems (the y-axis is in log scale).

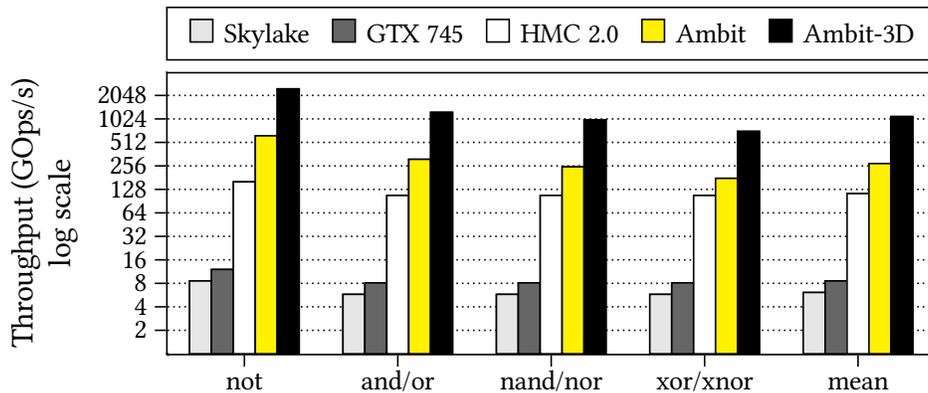


Figure 21: Throughput of bulk bitwise operations.

We draw three conclusions. First, the throughput of Skylake, GTX 745, and HMC 2.0 are limited by the memory bandwidth available to the respective processors. With an order of magnitude higher available memory bandwidth, HMC 2.0 achieves 18.5X and 13.1X better throughput for bulk bitwise operations compared to Skylake and GTX 745, respectively. Second, Ambit, with its ability to exploit the maximum internal DRAM bandwidth and memory-level parallelism, outperforms all three systems. On average, Ambit (with 8 DRAM banks) outperforms Skylake by 44.9X, GTX 745 by 32.0X, and HMC 2.0 by 2.4X. Third, 3D-stacked DRAM architectures like HMC contain a large number of banks (256 banks in 4GB HMC 2.0). By extending 3D-stacked DRAM with support for Ambit, Ambit-3D improves the throughput of bulk bitwise operations by 9.7X compared to HMC 2.0.

We estimate energy for DDR3-1333 using the Rambus power model [13]. Our energy numbers include only the DRAM and channel energy, and not the energy consumed by the processor. For Ambit, some activations have to raise multiple wordlines and hence, consume higher energy. Based on our analysis, the activation energy increases by 22% for each additional wordline raised. Table 4 shows the energy consumed per kilo-byte for different bitwise operations. Across all bitwise operations, Ambit reduces energy consumption by 25.1X—59.5X compared to copying data with the memory controller using the DDR3 interface.

	Design	not	and/or	nand/nor	xor/xnor
DRAM &	DDR3	93.7	137.9	137.9	137.9
Channel Energy	Ambit	1.6	3.2	4.0	5.5
(nJ/KB)	(↓)	59.5X	43.9X	35.1X	25.1X

Table 4: Energy of bitwise operations. (↓) indicates energy reduction of Ambit over the traditional DDR3-based design.

8. Effect on Real-World Applications

We evaluate the benefits of Ambit on real-world applications using the Gem5 full-system simulator [22]. Table 5 lists the main simulation parameters. Our simulations take into account the cost of maintaining coherence, and the overhead of RowClone to perform copy operations. We assume that application data is mapped such that all bitwise operations happen across rows *within a subarray*. We quantitatively evaluate three applications: 1) a database bitmap index [114, 10, 11, 5], 2) BitWeaving [97], a mechanism to accelerate database column scan operations, and 3) a bitvector-based implementation of the widely-used *set* data structure. In Section 8.4, we discuss four other applications that can benefit from Ambit.

Processor	x86, 8-wide, out-of-order, 4 Ghz 64-entry instruction queue
L1 cache	32 KB D-cache, 32 KB I-cache, LRU policy
L2 cache	2 MB, LRU policy, 64 B cache line size
Memory Controller	8 KB row size, FR-FCFS [127, 158] scheduling
Main memory	DDR4-2400, 1-channel, 1-rank, 16 banks

Table 5: Major simulation parameters

8.1. Bitmap Indices

Bitmap indices [28] are an alternative to traditional B-tree indices for databases. Compared to B-trees, bitmap indices 1) consume less space, and 2) can perform better for many queries (e.g., joins, scans). Several major databases support bitmap indices (e.g., Oracle [114], Redis [10], Fastbit [5], rlite [11]). Several real applications (e.g., [12, 2, 3, 36]) use bitmap indices for fast analytics. As bitmap indices heavily rely on bulk bitwise operations, Ambit can accelerate bitmap indices, thereby improving overall application performance.

To demonstrate this benefit, we use the following workload from a real application [36]. The application uses bitmap indices to track users’ characteristics (e.g., gender) and activities (e.g., did the user log in to the website on day ‘X’?) for u users. Our workload runs the following query: “How many unique users were active every week for the past w weeks? and How many male users were active each of the past w weeks?” Executing this query requires $6w$ bulk bitwise `or`, $2w-1$ bulk bitwise `and`, and $w+1$ bulk bitcount operations. In our mechanism, the bitcount operations are performed by the CPU. Figure 22 shows the end-to-end query execution time of the baseline and Ambit for the above experiment for various values of u and w .

We draw two conclusions. First, as each query has $O(w)$ bulk bitwise operations and each bulk bitwise operation takes $O(u)$ time, the query execution time increases with increasing value uw . Second, Ambit significantly reduces the query execution time compared to the baseline, by 6X on average.

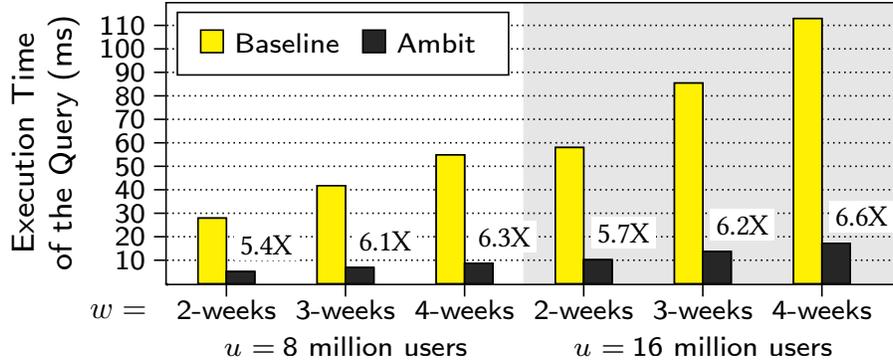


Figure 22: Bitmap index performance. The value above each bar indicates the reduction in execution time due to Ambit.

While we demonstrate the benefits of Ambit using one query, as all bitmap index queries involve several bulk bitwise operations, we expect Ambit to provide similar performance benefits for any application using bitmap indices.

8.2. BitWeaving: Fast Scans using Bitwise Operations

Column scan operations are a common part of many database queries. They are typically performed as part of evaluating a predicate. For a column with integer values, a predicate is typically of the form, $c1 \leq val \leq c2$, for two integer constants $c1$ and $c2$. Recent works [97, 148] observe that existing data representations for storing columnar data are inefficient for such predicate evaluation especially when the number of bits used to store each value of the column is less than the processor word size (typically 32 or 64). This is because 1) the values do not align well with word boundaries, and 2) the processor typically does not have comparison instructions at granularities smaller than the word size. To address this problem, BitWeaving [97] proposes two column representations, called BitWeaving-H and BitWeaving-V. As BitWeaving-V is faster than BitWeaving-H, we focus our attention on BitWeaving-V, and refer to it as just BitWeaving.

BitWeaving stores the values of a column such that the first bit of all the values of the column are stored contiguously, the second bit of all the values of the column are stored contiguously, and so on. Using this representation, the predicate $c1 \leq val \leq c2$, can be represented as a series of bitwise operations starting from the most significant bit all the way to the least significant bit (we refer the reader to the BitWeaving paper [97] for the detailed algorithm). As these bitwise operations can be performed in parallel across multiple values of the column, BitWeaving uses the hardware SIMD support to accelerate these operations. With support for Ambit, these operations can be performed in parallel across a larger set of values compared to 128/256-bit SIMD available in existing CPUs, thereby enabling higher performance.

We show this benefit by comparing the performance of BitWeaving using a baseline CPU with support for 128-bit SIMD to the performance of BitWeaving accelerated by Ambit for the following commonly-used query on a table T:

```
'select count(*) from T where c1 <= val <= c2'
```

Evaluating the predicate involves a series of bulk bitwise operations and the `count(*)` requires a bitcount operation. The execution time of the query depends on 1) the number of bits (b) used to represent each value `val`, and 2) the number of rows (r) in the table `T`. Figure 23 shows the speedup of Ambit over the baseline for various values of b and r .

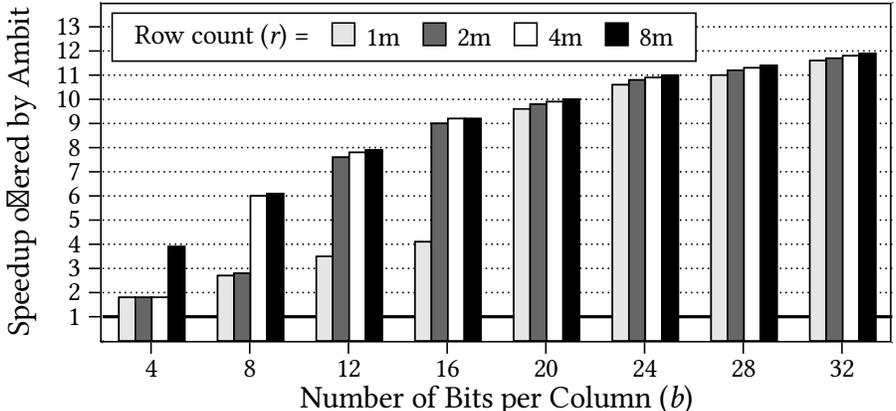


Figure 23: Speedup offered by Ambit for BitWeaving over our baseline CPU with SIMD support

We draw three conclusions. First, Ambit improves the performance of the query by between 1.8X and 11.8X (7.0X on average) compared to the baseline for various values of b and r . Second, the performance improvement of Ambit increases with increasing number of bits per column (b), because, as b increases, the fraction of time spent in performing the bitcount operation reduces. As a result, a larger fraction of the execution time can be accelerated using Ambit. Third, for $b = 4, 8, 12,$ and 16 , we observe large jumps in the speedup of Ambit as we increase the row count. These large jumps occur at points where the working set stops fitting in the on-chip cache. By exploiting the high bank-level parallelism in DRAM, Ambit can outperform the baseline (by up to 4.1X) *even when* the working set fits in the cache.

8.3. Bitvectors vs. Red-Black Trees

Many algorithms heavily use the *set* data structure. Red-black trees [48] (RB-trees) are typically used to implement a set [4]. However, a set with a limited domain can be implemented using a bitvector—a set that contains only elements from 1 to N , can be represented using an N -bit bitvector (e.g., Bitset [4]). Each bit indicates whether the corresponding element is present in the set. Bitvectors provide constant-time *insert* and *lookup* operations compared to $O(\log n)$ time taken by RB-trees. However, set operations like *union*, *intersection*, and *difference* have to scan the *entire* bitvector regardless of the number of elements actually present in the set. As a result, for these three operations, depending on the number of elements in the set, bitvectors may outperform or perform worse than RB-trees. With support for fast bulk bitwise operations, we show that Ambit significantly shifts the trade-off in favor of bitvectors for these three operations.

To demonstrate this, we compare the performance of set *union*, *intersection*, and *difference* using: RB-tree, bitvectors with 128-bit SIMD support (Bitset), and bitvectors with Ambit. We run a benchmark that performs each operation on m input sets and stores the result in an output set. We restrict the domain of the elements to be from 1 to N . Therefore, each set can be represented using an N -bit bitvector. For each of the three operations, we run multiple experiments varying the number of elements (e) *actually* present in each input set. Figure 24 shows the execution time of RB-tree, Bitset, and Ambit normalized to RB-tree for the three operations for $m = 15$, and $N = 512k$.

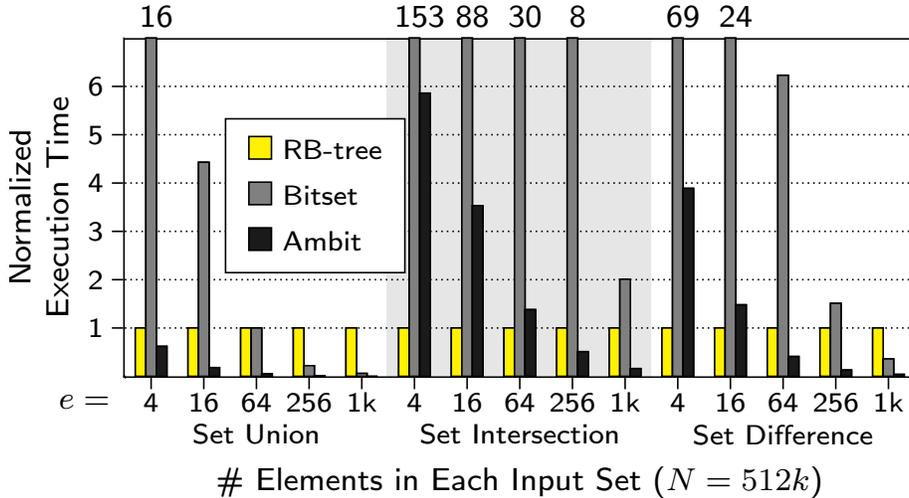


Figure 24: Performance of set operations

We draw three conclusions. First, by enabling much higher throughput for bulk bitwise operations, Ambit outperforms the baseline Bitset on all the experiments. Second, as expected, when the number of elements in each set is very small (16 out of $512k$), RB-Tree performs better than Bitset and Ambit (with the exception of *union*). Third, even when each set contains only 64 or more elements out of $512k$, Ambit significantly outperforms RB-Tree, 3X on average. We conclude that Ambit makes the bitvector-based implementation of a set more attractive than the commonly-used red-black-tree-based implementation.

8.4. Other Applications

We describe five other examples of applications that can significantly benefit from Ambit in terms of both performance and energy efficiency. We leave the evaluation of these applications with Ambit to future works.

8.4.1. BitFunnel: Web Search

Web search is an important workload in modern systems. From the time a query is issued to the time the results are sent back, web search involves several steps. Document filtering is one of the most time consuming steps. It identifies all documents that contain all the words in the input query. Microsoft recently open-sourced BitFunnel [47], a technology that

improves the efficiency of document filtering. BitFunnel represents both documents and queries as a bag of words using Bloom filters [23]. It then uses bitwise AND operations on specific locations of the Bloom filters to efficiently identify documents that contain all the query words. With Ambit, this operation can be significantly accelerated by simultaneously performing the filtering for thousands of documents.

8.4.2. Masked Initialization

Masked initializations [120] are very useful in applications like graphics (e.g., for clearing a specific color in an image). Masked operations can be represented using bitwise AND or OR operations. For example, the i^{th} bit of an integer x can be set using $x = x|(1 \ll i)$. Similarly, the i^{th} bit can be reset using $x = x\&(1 \ll i)$. These masks can be preloaded into DRAM rows and can be used to perform masked operations on a large amount of data. This operation can be easily accelerated using Ambit.

8.4.3. Encryption

Many encryption algorithms heavily use bitwise operations (e.g., XOR) [143, 51, 106]. Ambit’s support for fast bulk bitwise operations can i) boost the performance of existing encryption algorithms, and ii) enable new encryption algorithms with high throughput and efficiency.

8.4.4. DNA Sequence Mapping

In DNA sequence mapping, prior works [94, 128, 147, 151, 150, 123, 21, 18, 77, 89, 129, 17, 19] propose algorithms to map sequenced *reads* to the reference genome. Some works [150, 123, 21, 18, 77, 111, 17] heavily use bulk bitwise operations. Ambit can significantly improve the performance of such DNA sequence mapping algorithms [77].

8.4.5. Machine Learning

Recent works have shown that deep neural networks (DNNs) outperform other machine learning techniques in many problems such as image classification and speech recognition. Given the high computational requirements of DNNs, there has been a focus on neural networks with binary values [124, 60]. A recent work also exploits bit-serial computation to execute DNN inference algorithms using SIMD bitwise operations in the on-chip SRAM cache [38]. In conjunction with such techniques that increase the fraction of bitwise operations in DNNs, Ambit can significantly improve the performance of DNN algorithms.

9. Future Work

There are several avenues for future work that can build on top of the Ambit proposal. We briefly describe these avenues in this section.

9.1. Extending Ambit to Other Operations

We envision two major class of operations that can significantly increase the domain of applications that can benefit from Ambit.

The first operation is *count*. Counting the number of non-zero bits can be a very useful operation in many applications, including the ones evaluated in this paper (Section 8). Extending this operation to count the number of non-zero integers (of different widths) can further expand the scope of this operation.

The second operation is *shift*. Most arithmetic operations require some kind of bitwise shift. Shifting is also heavily used in any encryption algorithms [143, 51, 106] and bioinformatics algorithms [21, 150, 18, 77, 111, 17]. Extending Ambit to support bit shifting at different granularities can allow Ambit to significantly accelerate these applications. While recent works [95, 35] explore a possible implementation of shifting inside DRAM, several challenges remain unaddressed (e.g., data mapping, handling column failures in DRAM).

9.2. Evaluation of New Applications with Ambit

In Section 8.4, we discussed five other applications that can benefit significantly from Ambit. A concrete evaluation of these applications can 1) further strengthen the case for Ambit, and 2) reveal other potential operations that may require acceleration in DRAM in order to obtain good end-to-end application speedups (e.g., *bitcount*).

9.3. Redesigning Applications to Exploit Ambit

In Section 8.2, we show how the BitWeaving [97] technique can be accelerated using Ambit. BitWeaving is a technique that *redesigns* the data layout of tables in a database and uses an appropriate algorithm to execute large scan queries efficiently. It is the modified data layout that makes BitWeaving amenable for SIMD/Ambit acceleration. We believe similar careful redesign techniques can be used for other important workloads such as graph processing, machine learning, bioinformatics algorithms, etc., that can make them very amenable for acceleration using Ambit.

9.4. Taking Advantage of Approximate Ambit

Ambit requires a costly ECC mechanism to avoid errors during computation. In addition, the error rate may increase with increasing process variation. However, if an application can tolerate errors in computation, as observed by prior works [107, 104, 153, 142], then Ambit can be used as an approximate in-DRAM computation substrate.

10. Conclusion

In this paper, we focused on bulk bitwise operations, a class of operations heavily used by some important applications. Existing systems are inefficient at performing such operations

as they have to transfer a large amount of data on the memory channel, resulting in high latency, high memory bandwidth consumption, and high energy consumption.

We described Ambit, which employs the notion of *Processing using Memory* introduced in a recent work [136]. Ambit converts DRAM-based main memory into a bulk bitwise operation execution engine that performs bulk bitwise operations completely inside the memory [134]. Ambit has two component mechanisms. The first mechanism exploits the fact that many DRAM cells share the same sense amplifier and uses the idea of simultaneous activation of three rows of DRAM cells to perform bitwise MAJORITY/AND/OR operations efficiently. The second mechanism uses the inverters already present inside DRAM sense amplifiers to efficiently perform bitwise NOT operations. Since Ambit heavily exploits the internal organization and operation of DRAM, it incurs very low cost on top of the commodity DRAM architecture ($< 1\%$ chip area overhead).

Our evaluations show that Ambit enables between one to two order of magnitude improvement in raw throughput and energy consumption of bulk bitwise operations compared to existing DDR interfaces. We describe many real-world applications that can take advantage of Ambit. Our evaluations with three such applications show that Ambit can improve average performance of these applications compared to the baseline by between 3.0X-11.8X. Given its low cost and large performance improvements, we believe Ambit is a promising execution substrate that can accelerate applications that heavily use bulk bitwise operations. We hope future work will uncover even more potential in Ambit like execution.

Acknowledgments

We thank the reviewers of ISCA 2016/2017, MICRO 2016/2017, and HPCA 2017 for their valuable comments on various drafts of shorter versions of this work. We thank the members of the SAFARI group and PDL for their feedback. We acknowledge the generous support of our industrial partners, over the years, especially AliBaba, Google, Huawei, Intel, Microsoft, Nvidia, Samsung, Seagate, and VMWare. This work was supported in part by NSF, SRC, and the Intel Science and Technology Center for Cloud Computing. Some components of this work appeared in IEEE CAL [132], MICRO [134], ADCOM [136], and Seshadri's Ph.D. thesis [130].

References

- [1] 6th Generation Intel Core Processor Family Datasheet. <http://www.intel.com/content/www/us/en/processors/core/desktop-6th-gen-core-family-datasheet-vol-1.html>.
- [2] Belly card engineering. <https://tech.bellycard.com/>.
- [3] bitmapist: Powerful realtime analytics with Redis 2.6's bitmaps and Python. <https://github.com/Doist/bitmapist>.

- [4] C++ containers library, `std::set`. <http://en.cppreference.com/w/cpp/container/set>.
- [5] FastBit: An Efficient Compressed Bitmap Index Technology. <https://sdm.lbl.gov/fastbit/>.
- [6] GeForce GTX 745. <http://www.geforce.com/hardware/desktop-gpus/geforce-gtx-745-oem/specifications>.
- [7] High Bandwidth Memory DRAM. <http://www.jedec.org/standards-documents/docs/jesd235>.
- [8] Hybrid Memory Cube Specification 2.0. http://www.hybridmemorycube.org/files/SiteDownloads/HMC-30G-VSR_HMCC_Specification_Rev2.0_Public.pdf.
- [9] Predictive Technology Model. <http://ptm.asu.edu/>.
- [10] Redis - bitmaps. <http://redis.io/topics/data-types-intro#bitmaps>.
- [11] rlite: A Self-contained, Serverless, Zero-configuration, Transactional Redis-compatible Database Engine. <https://github.com/seppo0010/rlite>.
- [12] Spool. <http://www.getspool.com/>.
- [13] DRAM Power Model. <https://www.rambus.com/energy/>, 2010.
- [14] Junwhan Ahn, Sungpack Hong, Sungjoo Yoo, Onur Mutlu, and Kiyoun Choi. A Scalable Processing-in-memory Accelerator for Parallel Graph Processing. In *ISCA*, 2015.
- [15] Junwhan Ahn, Sungjoo Yoo, Onur Mutlu, and Kiyoun Choi. PIM-enabled Instructions: A Low-overhead, Locality-aware Processing-in-memory Architecture. In *ISCA*, 2015.
- [16] Berkin Akin, Franz Franchetti, and James C. Hoe. Data Reorganization in Memory Using 3D-stacked DRAM. In *ISCA*, 2015.
- [17] Mohammed Alser, Hasan Hassan, Akash Kumar, Onur Mutlu, and Can Alkan. Shouji: A Fast and Efficient Pre-alignment Filter for Sequence Alignment. *Bioinformatics*, 03 2019.
- [18] Mohammed Alser, Hasan Hassan, Hongyi Xin, Oğuz Ergin, Onur Mutlu, and Can Alkan. GateKeeper: A New Hardware Architecture for Accelerating Pre-Alignment in DNA Short Read Mapping. *Bioinformatics*, 2017.
- [19] Mohammed Alser, Onur Mutlu, and Can Alkan. Magnet: understanding and improving the accuracy of genome pre-alignment filtering. *arXiv preprint arXiv:1707.01631*, 2017.
- [20] Oreoluwa Babarinsa and Stratos Idreos. JAFAR: Near-Data Processing for Databases. In *SIGMOD*, 2015.

- [21] Gary Benson, Yozen Hernandez, and Joshua Loving. A Bit-Parallel, General Integer-Scoring Sequence Alignment Algorithm. In *CPM*, 2013.
- [22] Nathan Binkert, Bradford Beckmann, Gabriel Black, Steven K. Reinhardt, Ali Saidi, Arkaprava Basu, Joel Hestness, Derek R. Hower, Tushar Krishna, Somayeh Sardashti, Rathijit Sen, Korey Sewell, Muhammad Shoaib, Nilay Vaish, Mark D. Hill, and David A. Wood. The gem5 simulator. *SIGARCH Comput. Archit. News*, 39, 2011.
- [23] Burton H. Bloom. Space/time Trade-offs in Hash Coding with Allowable Errors. *ACM Communications*, 13, July 1970.
- [24] Amirali Boroumand, Saugata Ghose, Youngsok Kim, Rachata Ausavarungnirun, Eric Shiu, Rahul Thakur, Daehyun Kim, Aki Kuusela, Allan Knies, Parthasarathy Ranganathan, and Onur Mutlu. Google Workloads for Consumer Devices: Mitigating Data Movement Bottlenecks. In *ASPLOS*, 2018.
- [25] Amirali Boroumand, Saugata Ghose, Brandon Lucia, Kevin Hsieh, Krishna Malladi, Hongzhong Zheng, and Onur Mutlu. LazyPIM: An Efficient Cache Coherence Mechanism for Processing-in-Memory. *IEEE CAL*, 2016.
- [26] Amirali Boroumand, Saugata Ghose, Minesh Patel, Rachata Ausavarungnirun, Hasan Hassan, Brandon Lucia, Kevin Hsieh, Nastaran Hajinazar, Krishna T. Malladi, and Onur Mutlu. CoNDA: Enabling Efficient Near-Data Accelerator Communication by Optimizing Data Movement. In *ISCA*, 2019.
- [27] Amirali Boroumand, Saugata Ghose, Minesh Patel, Hasan Hassan, Brandon Lucia, Nastaran Hajinazar, Kevin Hsieh, Krishna T. Malladi, Hongzhong Zheng, and Onur Mutlu. LazyPIM: Efficient Support for Cache Coherence in Processing-in-Memory Architectures. *CoRR*, abs/1706.03162, 2017.
- [28] Chee-Yong Chan and Yannis E. Ioannidis. Bitmap index design and evaluation. In *SIGMOD*, 1998.
- [29] Kevin Chang. *Understanding Reduced-Voltage Operation in Modern DRAM Chips: Characterization, Analysis, and Mechanisms*. PhD thesis, Carnegie Mellon University, 2017.
- [30] Kevin K. Chang, Abhijith Kashyap, Hasan Hassan, Saugata Ghose, Kevin Hsieh, Donghyuk Lee, Tianshi Li, Gennady Pekhimenko, Samira Khan, and Onur Mutlu. Understanding Latency Variation in Modern DRAM Chips: Experimental Characterization, Analysis, and Optimization. In *SIGMETRICS*, 2016.
- [31] Kevin K Chang, Prashant J Nair, Donghyuk Lee, Saugata Ghose, Moinuddin K Qureshi, and Onur Mutlu. Low-Cost Inter-Linked Subarrays (LISA): Enabling Fast Inter-Subarray Data Movement in DRAM. In *HPCA*, 2016.

- [32] Kevin K Chang, A Giray Yaglikçi, Saugata Ghose, Aditya Agrawal, Niladrish Chatterjee, Abhijith Kashyap, Donghyuk Lee, Mike O’Connor, Hasan Hassan, and Onur Mutlu. Understanding Reduced-voltage Operation in Modern DRAM Devices: Experimental Characterization, Analysis, and Mechanisms. *SIGMETRICS*, 2017.
- [33] Kevin Kai-Wei Chang, Donghyuk Lee, Zeshan Chishti, Alaa R Alameldeen, Chris Wilkerson, Yoongu Kim, and Onur Mutlu. Improving DRAM performance by parallelizing refreshes with accesses. In *HPCA*, 2014.
- [34] J. Corbet, A. Rubini, and G. Kroah-Hartman. *Linux Device Drivers*, page 445. O’Reilly Media, 2005.
- [35] Quan Deng, Lei Jiang, Youtao Zhang, Minxuan Zhang, and Jun Yang. DrAcc: A DRAM Based Accelerator for Accurate CNN Inference. In *DAC*, 2018.
- [36] Demiz Denir, Islam AbdelRahman, Liang He, and Yingsheng Gao. Audience Insights query engine: In-memory integer store for social analytics. [https://code.facebook.com/posts/382299771946304/audience-insights-query-engine-in-memory-integer-store-for-social-analytics-/](https://code.facebook.com/posts/382299771946304/audience-insights-query-engine-in-memory-integer-store-for-social-analytics/).
- [37] Paul Dlugosch, Dave Brown, Paul Glendenning, Michael Leventhal, and Harold Noyes. An Efficient and Scalable Semiconductor Architecture for Parallel Automata Processing. *IEEE TPDS*, 2014.
- [38] Charles Eckert, Xiaowei Wang, Jingcheng Wang, Arun Subramaniyan, Ravi Iyer, Dennis Sylvester, David Blaauw, and Reetuparna Das. Neural cache: Bit-serial in-cache acceleration of deep neural networks. In *ISCA*, 2018.
- [39] Duncan Elliott, Michael Stumm, W. Martin Snelgrove, Christian Cojocar, and Robert McKenzie. Computational RAM: Implementing Processors in Memory. *IEEE Des. Test*, 16, 1999.
- [40] C. F. Falconer, C. P. Mozak, and A. J. Normal. Suppressing Power Supply Noise Using Data Scrambling in Double Data Rate Memory Systems. US Patent 8503678, 2009.
- [41] A. Farmahini-Farahani, Jung Ho Ahn, K. Morrow, and Nam Sung Kim. NDA: Near-DRAM acceleration architecture leveraging commodity DRAM devices and standard memory modules. In *HPCA*, 2015.
- [42] Mingyu Gao, Grant Ayers, and Christos Kozyrakis. Practical Near-Data Processing for In-Memory Analytics Frameworks. In *PACT*, 2015.
- [43] Mingyu Gao and Christos Kozyrakis. HRL: Efficient and Flexible Reconfigurable Logic for Near-Data Processing. In *HPCA*, 2016.
- [44] Saugata Ghose, Abdullah Giray Yaglikçi, Raghav Gupta, Donghyuk Lee, Kais Kudrolli, William X. Liu, Hasan Hassan, Kevin K. Chang, Niladrish Chatterjee, Aditya Agrawal, Mike O’Connor, and Onur Mutlu. What Your DRAM Power Models Are Not Telling You: Lessons from a Detailed Experimental Study. In *SIGMETRICS*, 2018.

- [45] Andrew Glew. MLP yes! ILP no. *ASPLOS WACI*, 1998.
- [46] Maya Gokhale, Bill Holmes, and Ken Iobst. Processing in Memory: The Terasys Massively Parallel PIM Array. *Computer*, 28, 1995.
- [47] Bob Goodwin, Michael Hopcroft, Dan Luu, Alex Clemmer, Mihaela Curmei, Sameh Elnikety, and Yuxiong He. BitFunnel: Revisiting Signatures for Search. In *SIGIR*, 2017.
- [48] Leo J. Guibas and Robert Sedgwick. A Dichromatic Framework for Balanced Trees. In *SFCS*, 1978.
- [49] Qi Guo, Nikolaos Alachiotis, Berkin Akin, Fazle Sadi, Guanglin Xu, Tze Meng Low, Larry Pillegi, James C. Hoe, and Franz Frachetti. 3D-Stacked Memory-Side Acceleration: Accelerator and System Design. In *WoNDP*, 2013.
- [50] R. W. Hamming. Error Detecting and Error Correcting Codes. *BSTJ*, 1950.
- [51] Jong-Wook Han, Choon-Sik Park, Dae-Hyun Ryu, and Eun-Soo Kim. Optical image encryption based on XOR operations. *Optical Engineering*, 38, 1999.
- [52] Milad Hashemi, Khubaib, Eiman Ebrahimi, Onur Mutlu, and Yale N. Patt. Accelerating Dependent Cache Misses with an Enhanced Memory Controller. In *ISCA*, 2016.
- [53] Milad Hashemi, Onur Mutlu, and Yale N. Patt. Continuous Runahead: Transparent Hardware Acceleration for Memory Intensive Workloads. In *MICRO*, 2016.
- [54] Hasan Hassan, Minesh Patel, Jeremie Kim, A. Giray Yaglikçi, Nandita Vijaykumar, Nika Mansouri Ghiasi, Saugata Ghose, and Onur Mutlu. CROW: A Low-Overhead Substrate for Improving DRAM Performance and Energy-Efficiency. In *ISCA*, 2019.
- [55] Hasan Hassan, Gennady Pekhimenko, Nandita Vijaykumar, Vivek Seshadri, Donghyuk Lee, Oguz Ergin, and Onur Mutlu. ChargeCache: Reducing DRAM Latency by Exploiting Row Access Locality. In *HPCA*, 2016.
- [56] Hasan Hassan, Nandita Vijaykumar, Samira Khan, Saugata Ghose, Kevin Chang, Gennady Pekhimenko, Donghyuk Lee, Oguz Ergin, and Onur Mutlu. SoftMC: A Flexible and Practical Open-source Infrastructure for Enabling Experimental DRAM Studies. In *HPCA*, 2017.
- [57] Syed Minhaj Hassan, Sudhakar Yalamanchili, and Saibal Mukhopadhyay. Near Data Processing: Impact and Optimization of 3D Memory System Architecture on the Uncore. In *MEMSYS*, 2015.
- [58] Kevin Hsieh, Eiman Ebrahimi, Gwangsun Kim, Niladrish Chatterjee, Mike O’Conner, Nandita Vijaykumar, Onur Mutlu, and Stephen W. Keckler. Transparent Offloading and Mapping (TOM): Enabling Programmer-Transparent Near-Data Processing in GPU Systems. In *ISCA*, 2016.

- [59] Kevin Hsieh, Samira Khan, Nandita Vijaykumar, Kevin K Chang, Amirali Boroumand, Saugata Ghose, and Onur Mutlu. Accelerating Pointer Chasing in 3D-stacked Memory: Challenges, Mechanisms, Evaluation. In *ICCD*, 2016.
- [60] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized Neural Networks. In *NIPS*, 2016.
- [61] Intel. Intel Instruction Set Architecture Extensions. <https://software.intel.com/en-us/intel-isa-extensions>.
- [62] Engin Ipek, Onur Mutlu, José F. Martínez, and Rich Caruana. Self-Optimizing Memory Controllers: A Reinforcement Learning Approach. In *ISCA*, ISCA '08, 2008.
- [63] J. Jeddelloh and B. Keeth. Hybrid Memory Cube: New DRAM architecture increases density and performance. In *VLSIT*, 2012.
- [64] JEDEC. DDR3 SDRAM Standard, JESD79-3D. <http://www.jedec.org/sites/default/files/docs/JESD79-3D.pdf>, 2009.
- [65] JEDEC. DDR3 SDRAM, JESD79-3F, 2012.
- [66] JEDEC. DDR4 SDRAM Standard. <http://www.jedec.org/standards-documents/docs/jesd79-4a>, 2013.
- [67] H.B. Kang and S.K. Hong. One-Transistor Type DRAM. US Patent 7701751, 2009.
- [68] Uksong Kang, Hak-soo Yu, Churoo Park, Hongzhong Zheng, John Halbert, Kuljit Bains, S Jang, and Joo Sun Choi. Co-architecting Controllers and DRAM to Enhance DRAM Process Scaling. In *The Memory Forum*, 2014.
- [69] Brent Keeth, R. Jacob Baker, Brian Johnson, and Feng Lin. *DRAM Circuit Design: Fundamental and High-Speed Topics*. Wiley-IEEE Press, 2nd edition, 2007.
- [70] S. Khan, C. Wilkerson, D. Lee, A. R. Alameldeen, and O. Mutlu. A Case for Memory Content-Based Detection and Mitigation of Data-Dependent Failures in DRAM. *IEEE CAL*, 2017.
- [71] Samira Khan, Donghyuk Lee, Yoongu Kim, Alaa R. Alameldeen, Chris Wilkerson, and Onur Mutlu. The Efficacy of Error Mitigation Techniques for DRAM Retention Failures: A Comparative Experimental Study. In *SIGMETRICS*, 2014.
- [72] Samira Khan, Chris Wilkerson, Zhe Wang, Alaa R. Alameldeen, Donghyuk Lee, and Onur Mutlu. Detecting and Mitigating Data-dependent DRAM Failures by Exploiting Current Memory Content. In *MICRO*, 2017.
- [73] Samira M. Khan, Donghyuk Lee, and Onur Mutlu. PARBOR: An Efficient System-Level Technique to Detect Data-Dependent Failures in DRAM. In *DSN*, 2016.
- [74] Jeremie Kim, Minesh Patel, Hasan Hassan, and Onur Mutlu. Solar-DRAM: Reducing DRAM Access Latency by Exploiting the Variation in Local Bitlines. In *ICCD*, 2018.

- [75] Jeremie Kim, Minesh Patel, Hasan Hassan, and Onur Mutlu. The DRAM Latency PUF: Quickly Evaluating Physical Unclonable Functions by Exploiting the Latency-Reliability Tradeoff in Modern Commodity DRAM Devices. In *HPCA*, 2018.
- [76] Jeremie Kim, Minesh Patel, Hasan Hassan, Lois Orosa, and Onur Mutlu. D-RaNGe: Using Commodity DRAM Devices to Generate True Random Numbers with Low Latency and High Throughput. In *HPCA*, 2019.
- [77] Jeremie S Kim, Damla Senol Cali, Hongyi Xin, Donghyuk Lee, Saugata Ghose, Mohammed Alser, Hasan Hassan, Oguz Ergin, Can Alkan, and Onur Mutlu. GRIM-Filter: Fast seed location filtering in DNA read mapping using processing-in-memory technologies. *BMC genomics*, 19, 2018.
- [78] Jeremie S. Kim, Minesh Patel, Hasan Hassan, and Onur Mutlu. The DRAM Latency PUF: Quickly Evaluating Physical Unclonable Functions by Exploiting the Latency-Reliability Tradeoff in Modern Commodity DRAM Devices. In *HPCA*, 2018.
- [79] Yoongu Kim. *Architectural Techniques to Enhance DRAM Scaling*. PhD thesis, Carnegie Mellon University, 2015.
- [80] Yoongu Kim, Ross Daly, Jeremie Kim, Chris Fallin, Ji Hye Lee, Donghyuk Lee, Chris Wilkerson, Konrad Lai, and Onur Mutlu. Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors. In *ISCA*, 2014.
- [81] Yoongu Kim, Michael Papamichael, Onur Mutlu, and Mor Harchol-Balter. Thread Cluster Memory Scheduling: Exploiting Differences in Memory Access Behavior. In *MICRO*, 2010.
- [82] Yoongu Kim, Vivek Seshadri, Donghyuk Lee, Jamie Liu, and Onur Mutlu. A Case for Exploiting Subarray-level Parallelism (SALP) in DRAM. In *ISCA*, 2012.
- [83] Yoongu Kim, Weikun Yang, and Onur Mutlu. Ramulator: A Fast and Extensible DRAM Simulator. *IEEE CAL*, 2016.
- [84] D. E. Knuth. *The Art of Computer Programming. Fascicle 1: Bitwise Tricks & Techniques; Binary Decision Diagrams*, 2009.
- [85] Peter M. Kogge. EXECUBE: A New Architecture for Scaleable MPPs. In *ICPP*, 1994.
- [86] C. J. Lee, V. Narasiman, O. Mutlu, and Y. N. Patt. Improving Memory Bank-Level Parallelism in the Presence of Prefetching. In *MICRO*, 2009.
- [87] Donghyuk Lee. *Reducing DRAM Latency at Low Cost by Exploiting Heterogeneity*. PhD thesis, Carnegie Mellon University, 2016.
- [88] Donghyuk Lee, Saugata Ghose, Gennady Pekhimenko, Samira Khan, and Onur Mutlu. Simultaneous Multi-Layer Access: Improving 3D-Stacked Memory Bandwidth at Low Cost. *ACM TACO*, 12, 2016.

- [89] Donghyuk Lee, Farhad Hormozdiari, Hongyi Xin, Faraz Hach, Onur Mutlu, and Can Alkan. Fast and Accurate Mapping of Complete Genomics Reads. *Methods*, 2015.
- [90] Donghyuk Lee, Samira Khan, Lavanya Subramanian, Saugata Ghose, Rachata Ausavarungnirun, Gennady Pekhimenko, Vivek Seshadri, and Onur Mutlu. Design-Induced Latency Variation in Modern DRAM Chips: Characterization, Analysis, and Latency Reduction Mechanisms. In *SIGMETRICS*, 2017.
- [91] Donghyuk Lee, Yoongu Kim, Gennady Pekhimenko, Samira Manabi Khan, Vivek Seshadri, Kevin Kai-Wei Chang, and Onur Mutlu. Adaptive-latency DRAM: Optimizing DRAM timing for the common-case. In *HPCA*, 2015.
- [92] Donghyuk Lee, Yoongu Kim, Vivek Seshadri, Jamie Liu, Lavanya Subramanian, and Onur Mutlu. Tiered-latency DRAM: A Low Latency and Low Cost DRAM Architecture. In *HPCA*, 2013.
- [93] Donghyuk Lee, Lavanya Subramanian, Rachata Ausavarungnirun, Jongmoo Choi, and Onur Mutlu. Decoupled Direct Memory Access: Isolating CPU and IO Traffic by Leveraging a Dual-Data-Port DRAM. In *PACT*, 2015.
- [94] Heng Li and Richard Durbin. Fast and accurate long-read alignment with Burrows–Wheeler transform. *Bioinformatics*, 26, 2010.
- [95] Shuangchen Li, Dimin Niu, Krishna T. Malladi, Hongzhong Zheng, Bob Brennan, and Yuan Xie. DRISA: A DRAM-based Reconfigurable In-Situ Accelerator. In *MICRO*, 2017.
- [96] Shuangchen Li, Cong Xu, Qiaosha Zou, Jishen Zhao, Yu Lu, and Yuan Xie. Pinatubo: A Processing-in-Memory Architecture for Bulk Bitwise Operations in Emerging Non-Volatile Memories. In *DAC*, 2016.
- [97] Yinan Li and Jignesh M. Patel. BitWeaving: Fast Scans for Main Memory Data Processing. In *SIGMOD*, 2013.
- [98] Yinan Li and Jignesh M. Patel. WideTable: An Accelerator for Analytical Data Processing. *Proc. VLDB Endow.*, 2014.
- [99] Erik Lindholm, John Nickolls, Stuart Oberman, and John Montrym. NVIDIA Tesla: A Unified Graphics and Computing Architecture. *IEEE Micro*, 2008.
- [100] Jamie Liu, Ben Jaiyen, Yoongu Kim, Chris Wilkerson, and Onur Mutlu. An Experimental Study of Data Retention Behavior in Modern DRAM Devices: Implications for Retention Time Profiling Mechanisms. In *ISCA*, 2013.
- [101] Jamie Liu, Ben Jaiyen, Richard Veras, and Onur Mutlu. RAIDR: Retention-Aware Intelligent DRAM Refresh. In *ISCA*, 2012.
- [102] Zhiyu Liu, Irina Calciu, Maurice Herlihy, and Onur Mutlu. Concurrent Data Structures for Near-Memory Computing. In *SPAA*, 2017.

- [103] Shih-Lien Lu, Ying-Chen Lin, and Chia-Lin Yang. Improving DRAM Latency with Dynamic Asymmetric Subarray. In *MICRO*, 2015.
- [104] Yixin Luo, Sriram Govindan, Bikash Sharma, Mark Santaniello, Justin Meza, Aman Kansal, Jie Liu, Badriddine Khessib, Kushagra Vaid, and Onur Mutlu. Characterizing Application Memory Error Vulnerability to Optimize Datacenter Cost via Heterogeneous-Reliability Memory. In *DSN*, 2014.
- [105] R. E. Lyons and W. Vanderkulk. The Use of Triple-Modular Redundancy to Improve Computer Reliability. *IBM JRD*, 1962.
- [106] S. A. Manavski. CUDA Compatible GPU as an Efficient Hardware Accelerator for AES Cryptography. In *ICSPC*, 2007.
- [107] Sparsh Mittal. A Survey of Techniques for Approximate Computing. *ACM Comput. Surv.*, 48, 2016.
- [108] Amir Morad, Leonid Yavits, and Ran Ginosar. GP-SIMD Processing-in-Memory. *ACM Trans. Archit. Code Optim.*, 11, 2015.
- [109] Onur Mutlu and Thomas Moscibroda. Parallelism-Aware Batch Scheduling: Enhancing Both Performance and Fairness of Shared DRAM Systems. In *ISCA*, 2008.
- [110] Onur Mutlu, Jared Stark, Chris Wilkerson, and Yale N. Patt. Runahead Execution: An Alternative to Very Large Instruction Windows for Out-of-Order Processors. In *HPCA*, 2003.
- [111] Gene Myers. A Fast Bit-vector Algorithm for Approximate String Matching Based on Dynamic Programming. *JACM*, 1999.
- [112] L. Nai, R. Hadidi, J. Sim, H. Kim, P. Kumar, and H. Kim. GraphPIM: Enabling Instruction-Level PIM Offloading in Graph Computing Frameworks. In *HPCA*, 2017.
- [113] Elizabeth O’Neil, Patrick O’Neil, and Kesheng Wu. Bitmap Index Design Choices and Their Performance Implications. In *IDEAS*, 2007.
- [114] Oracle. Using Bitmap Indexes in Data Warehouses. https://docs.oracle.com/cd/B28359_01/server.111/b28313/indexe\s.htm.
- [115] Mark Oskin, Frederic T. Chong, and Timothy Sherwood. Active Pages: A Computation Model for Intelligent Memory. In *ISCA*, 1998.
- [116] Minesh Patel, Jeremie S. Kim, Hasan Hassan, and Onur Mutlu. Understanding and Modeling On-Die Error Correction in Modern DRAM: An Experimental Study Using Real Devices. In *DSN*, 2019.
- [117] Minesh Patel, Jeremie S Kim, and Onur Mutlu. The Reach Profiler (REAPER): Enabling the Mitigation of DRAM Retention Failures via Profiling at Aggressive Conditions. In *ISCA*, 2017.

- [118] David Patterson, Thomas Anderson, Neal Cardwell, Richard Fromm, Kimberly Keeton, Christoforos Kozyrakis, Randi Thomas, and Katherine Yelick. A Case for Intelligent RAM. *IEEE Micro*, 17, 1997.
- [119] Ashutosh Pattnaik, Xulong Tang, Adwait Jog, Onur Kayiran, Asit K Mishra, Mahmut T Kandemir, Onur Mutlu, and Chita R Das. SchEduling Techniques for GPU Architectures with Processing-in-memory Capabilities. In *PACT*, 2016.
- [120] Alex Peleg and Uri Weiser. MMX Technology Extension to the Intel Architecture. *IEEE Micro*, 1996.
- [121] M. K. Qureshi, D. H. Kim, S. Khan, P. J. Nair, and O. Mutlu. AVATAR: A Variable-Retention-Time (VRT) Aware Refresh for DRAM Systems. In *DSN*, 2015.
- [122] Rambus. DRAM power model, 2010.
- [123] Kim R Rasmussen, Jens Stoye, and Eugene W Myers. Efficient q-gram filters for finding all ε -matches over a given length. *Journal of Computational Biology*, 13(2), 2006.
- [124] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. XNOR-Net: ImageNet Classification Using Binary Convolutional Neural Networks. In *ECCV*, 2016.
- [125] P. J. Restle, J. W. Park, and B. F. Lloyd. DRAM Variable Retention Time. In *IEDM*, 1992.
- [126] Ronald L Rivest, Len Adleman, and Michael L Dertouzos. On Data Banks and Privacy Homomorphisms. *FSC*, 1978.
- [127] Scott Rixner, William J. Dally, Ujval J. Kapasi, Peter Mattson, and John D. Owens. Memory access scheduling. In *ISCA*, 2000.
- [128] Stephen M Rumble, Phil Lacroute, Adrian V Dalca, Marc Fiume, Arend Sidow, and Michael Brudno. SHRiMP: Accurate mapping of short color-space reads. 2009.
- [129] Damla Senol Cali, Jeremie S Kim, Saugata Ghose, Can Alkan, and Onur Mutlu. Nanopore Sequencing Technology and Tools for Genome Assembly: Computational Analysis of the Current State, Bottlenecks and Future Directions. *Briefings in Bioinformatics*, 2018.
- [130] Vivek Seshadri. *Simple DRAM and Virtual Memory Abstractions to Enable Highly Efficient Memory Subsystems*. PhD thesis, Carnegie Mellon University, 2016.
- [131] Vivek Seshadri, Abhishek Bhowmick, Onur Mutlu, Phillip B. Gibbons, Michael A. Kozuch, and Todd C. Mowry. The Dirty-Block Index. In *ISCA*, 2014.
- [132] Vivek Seshadri, Kevin Hsieh, Amirali Boroumand, Donghyuk Lee, Michael A. Kozuch, Onur Mutlu, Phillip B. Gibbons, and Todd C. Mowry. Fast Bulk Bitwise AND and OR in DRAM. *IEEE CAL*, 2015.

- [133] Vivek Seshadri, Yoongu Kim, Chris Fallin, Donghyuk Lee, Rachata Ausavarungnirun, Gennady Pekhimenko, Yixin Luo, Onur Mutlu, Phillip B. Gibbons, Michael A. Kozuch, and Todd C. Mowry. RowClone: Fast and Energy-efficient in-DRAM Bulk Data Copy and Initialization. In *MICRO*, 2013.
- [134] Vivek Seshadri, Donghyuk Lee, Thomas Mullins, Hasan Hassan, Amirali Boroumand, Jeremie Kim, Michael A. Kozuch, Onur Mutlu, Phillip B. Gibbons, and Todd C. Mowry. Ambit: In-Memory Accelerator for Bulk Bitwise Operations Using Commodity DRAM Technology. In *MICRO*, 2017.
- [135] Vivek Seshadri, Thomas Mullins, Amirali Boroumand, Onur Mutlu, Phillip B. Gibbons, Michael A. Kozuch, and Todd C. Mowry. Gather-scatter DRAM: In-DRAM Address Translation to Improve the Spatial Locality of Non-unit Strided Accesses. In *MICRO*, 2015.
- [136] Vivek Seshadri and Onur Mutlu. Simple Operations in Memory to Reduce Data Movement. In *Advances in Computers*, volume 106. Elsevier, 2017.
- [137] David Elliot Shaw, Salvatore Stolfo, Hussein Ibrahim, Bruce K. Hillyer, Jim Andrews, and Gio Wiederhold. The NON-VON Database Machine: An Overview. <http://hdl.handle.net/10022/AC:P:11530>., 1981.
- [138] Roman Sikorski. *Boolean Algebras*, volume 2. Springer, 1969.
- [139] Harold S. Stone. A Logic-in-Memory Computer. *IEEE Trans. Comput.*, 19, 1970.
- [140] Arun Subramaniyan and Reetuparna Das. Parallel Automata Processor. In *ISCA*, 2017.
- [141] Zehra Sura, Arpith Jacob, Tong Chen, Bryan Rosenburg, Olivier Sallenave, Carlo Bertolli, Samuel Antao, Jose Brunheroto, Yoonho Park, Kevin O'Brien, and Ravi Nair. Data access optimization in a processing-in-memory system. In *CF*, 2015.
- [142] Bradley Thwaites, Gennady Pekhimenko, Hadi Esmaeilzadeh, Amir Yazdanbakhsh, Onur Mutlu, Jongse Park, Girish Mururu, and Todd Mowry. Rollback-free Value Prediction with Approximate Loads. In *PACT*, 2014.
- [143] P. Tuyls, H. D. L. Hollmann, J. H. Van Lint, and L. Tolhuizen. XOR-based Visual Cryptography Schemes. *Designs, Codes and Cryptography*, 37, 2005.
- [144] Aniruddha N. Udipi, Naveen Muralimanohar, Niladrish Chatterjee, Rajeev Balasubramonian, Al Davis, and Norman P. Jouppi. Rethinking DRAM Design and Organization for Energy-constrained Multi-cores. In *ISCA*, 2010.
- [145] Yaohua Wang, Arash Tavakkol, Lois Orosa, Saugata Ghose, Nika Mansouri Ghiasi, Minesh Patel, Jeremie Kim, Hasan Hassan, Mohammad Sadrosadati, and Onur Mutlu. Reducing DRAM Latency via Charge-Level-Aware Look-Ahead Partial Restoration. In *MICRO*, 2018.

- [146] Henry S. Warren. *Hacker's Delight*. Addison-Wesley Professional, 2nd edition, 2012.
- [147] David Weese, Anne-Katrin Emde, Tobias Rausch, Andreas Döring, and Knut Reinert. Razers - fast read mapping with sensitivity control. *Genome research*, 19, 2009.
- [148] Thomas Willhalm, Ismail Oukid, Ingo Muller, and Franz Faerber. Vectorizing Database Column Scans with Complex Predicates. In *ADMS*, 2013.
- [149] Kesheng Wu, Ekow J. Otoo, and Arie Shoshani. Compressing Bitmap Indexes for Faster Search Operations. In *SSDBM*, 2002.
- [150] Hongyi Xin, John Greth, John Emmons, Gennady Pekhimenko, Carl Kingsford, Can Alkan, and Onur Mutlu. Shifted Hamming Distance: A Fast and Accurate SIMD-friendly Filter to Accelerate Alignment Verification in Read Mapping. *Bioinformatics*, 2015.
- [151] Hongyi Xin, Donghyuk Lee, Farhad Hormozdiari, Samihan Yedkar, Onur Mutlu, and Can Alkan. Accelerating read mapping with FastHASH. *BMC genomics*, 14:S13, 2013.
- [152] D. S. Yaney, C. Y. Lu, R. A. Kohler, M. J. Kelly, and J. T. Nelson. A Meta-stable Leakage Phenomenon in DRAM Charge Storage - Variable Hold Time. In *IEDM*, 1987.
- [153] Amir Yazdanbakhsh, Gennady Pekhimenko, Bradley Thwaites, Hadi Esmaeilzadeh, Onur Mutlu, and Todd C. Mowry. RFVP: Rollback-Free Value Prediction with Safe-to-Approximate Loads. *ACM TACO*, 2016.
- [154] Dongping Zhang, Nuwan Jayasena, Alexander Lyashevsky, Joseph L. Greathouse, Lifan Xu, and Michael Ignatowski. TOP-PIM: Throughput-oriented Programmable Processing in Memory. In *HPDC*, 2014.
- [155] Tao Zhang, Ke Chen, Cong Xu, Guangyu Sun, Tao Wang, and Yuan Xie. Half-DRAM: A High-bandwidth and Low-power DRAM Architecture from the Rethinking of Fine-grained Activation. In *ISCA*, 2014.
- [156] Wei Zhao and Yu Cao. New Generation of Predictive Technology Model for Sub-45 nm Early Design Exploration. *IEEE TED*, 2006.
- [157] Q. Zhu, T. Graf, H. E. Sumbul, L. Pileggi, and F. Franchetti. Accelerating sparse matrix-matrix multiplication with 3D-stacked logic-in-memory hardware. In *HPEC*, 2013.
- [158] W. K. Zuravleff and T. Robinson. Controller for a synchronous DRAM that maximizes throughput by allowing memory requests and commands to be issued out of order. Patent 5630096, 1997.