

CONCEPT: A Column Oriented Memory Controller for Efficient Memory and PIM Operations in RRAM

Nishil Talati*

University of Michigan,
Ann Arbor, Technion -
Israel Institute of
Technology

Heonjae Ha

Stanford University

Ben Perach

Technion - Israel Institute of
Technology

Ronny Ronen

Technion - Israel Institute of
Technology

Shahar Kvatinsky

Technion - Israel Institute of
Technology

While DRAM cannot easily scale below a 20nm technology node, RRAM suffers far less from scalability issues. Moreover, RRAM's resistivity enables its use for processing-in-memory (PIM), potentially alleviating the von Neumann bottleneck. Unfortunately, because of technological idiosyncrasies, existing DRAM-centric memory controllers cannot exploit the full potential of RRAM. In this paper, we present the design of a memory controller called CONCEPT. The controller is optimized to exploit unique properties of RRAM to enhance its performance and energy efficiency as well as exploiting RRAM's PIM capability. We show that with CONCEPT, RRAM can achieve DRAM-like performance and energy efficiency on SPEC CPU 2006 benchmarks. Furthermore, using RRAM PIM

capabilities, we show a 5X performance gain on a data-intensive in-memory database workload compared to a state-of-the-art CPU-memory computing model.

Index Terms—von Neumann bottleneck, memristor, Processing-In-Memory (PIM), Resistive RAM (RRAM), memory controller, memristor aided logic (MAGIC).

The process technology scaling of DRAM has so far facilitated low cost-per-bit by enabling reduction in cell size. However, further scaling has proven to be costly due to physical factors such as the difficulty of fabricating capacitors with high aspect ratio. Furthermore, because commodity DRAM cannot process data, it needs to be transferred to the CPU for processing via a bandwidth-limited bus. This data transfer severely limits the performance and the energy efficiency of computers¹, a limitation known as *the von Neumann bottleneck*.

The memory controller is a key component in optimizing the performance and energy efficiency of modern computing systems. Its purpose is to deliver performance and power optimizations by intelligently scheduling memory requests and accessing data given the limitations of a memory technology. DRAM has been the technological choice for building main memory systems for several decades mainly because of its low access latency and low cost. Therefore, existing memory controllers target DRAM, and their design constraints are focused around that specific technology.

*This work was carried out when the author was at the Technion – Israel Institute of Technology.

In this article, we propose to employ Resistive RAM (RRAM) to tackle both the DRAM scaling problem and the von Neumann bottleneck. While prior art employs a traditional DRAM memory controller to leverage unique properties of RRAM⁷, we propose a more fundamental and novel solution to the problem: *redesigning the memory controller to optimize the data access pattern as well as to perform processing-in-memory (PIM) using RRAM*.

We make several key observations about the technological idiosyncrasies that make DRAM and RRAM incompatible. Motivated by these observations, we propose a unique column-oriented access protocol for RRAM, which we call R-DDR, *i.e.*, an RRAM-access protocol similar to DDRx. This protocol is tailor-made around the design constraints of RRAM to best optimize its performance and energy efficiency. Optimizations proposed in the past⁷ are orthogonal to our approach which can be augmented on top of our approach, which may lead to further benefits in performance and energy efficiency of RRAM. In addition to supporting efficient memory operations, R-DDR can also support PIM operations using a technique called Memristor Aided loGIC (MAGIC)². MAGIC based PIM allows reducing data transfer to the CPU and leveraging the ample parallelism of the memory crossbar array³ to realize efficient single instruction multiple data (SIMD) operations. We extend the instruction set architecture (ISA) and hardware of the memory controller to support two additional instructions for PIM – two-input MAGIC NOR and MAGIC NOT. Note that our solution is motivated from how DDRx is highly optimized for a technology (*i.e.*, DRAM), however, we believe that in the future with heterogeneous memory system, programmable controller might be useful to support various protocols optimized for different memory technologies.

To evaluate our proposed system, we develop a comprehensive device-to-architecture simulation framework. Using this simulation framework, we compare performance and energy efficiency of our memory system against RRAM accessed using DDR4 on SPEC CPU2006 workloads. The results show that our approach improves performance by 35% and reduces energy by 7%. These improvements are the result of carefully designing the memory controller around the constraints of RRAM to optimize its memory access pattern. A comparison of our approach with other PIM approaches shows that our system outperforms PCM-based Pinatubo⁵ by 5X and achieves similar performance to DRAM-based Ambit⁴.

BACKGROUND AND MOTIVATION

DRAM Scaling Problem

DRAM uses a capacitor to store charge that is accessed by an access transistor. This capacitive element is generally fabricated in a vertically cylindrical structure to save chip area. To maintain the retention time requirement of the DRAM device, the capacitance needs to be constant at approximately 25 fF. DRAM scaling has so far been facilitated by scaling a linear dimension by 0.71X from generation to generation, and hence, the surface area of the capacitor scales by 0.5X. To compensate for the reduction in surface area, the height of the capacitor needs to be scaled by 2X to keep the capacitance value the same. The aspect ratio of the capacitor, which is the ratio of the height and its base diameter, increases exponentially with technology scaling, and goes above 100X at the sub-20nm technology node, hence, DRAM companies use costly sophisticated tricks to get around this issue. These narrow cylindrical cells are inherently unstable and mechanically difficult to fabricate. On the other hand, if the capacitor height is not scaled to compensate for the increase in area, the capacitor can store comparatively less charge. As a result, DRAM refresh rate must be increased, reducing the performance and increasing the energy consumption of the memory system. Therefore, the scaling of DRAM is a major challenge below the 20nm technology node.

The von Neumann Bottleneck

Contemporary computing systems are based on the von Neumann architecture, or an improvised version of it, where data is processed within the CPU and stored in memory (DRAM being main memory). Since traditional DRAM cells are charge-based and they cannot process data using the same cells that are used for storing data, data must be transferred outside the memory bank or chip for processing. The resulting data movement between the memory and the CPU causes major bottlenecks in terms of performance and energy efficiency¹. A few contemporary approaches propose to use DRAM sense amplifiers for computation. However, doing so would still require data movement from the memory cells to the sense amplifiers and back; hence, they are limited by the bandwidth available for data transfer between cell array and peripheral circuits. One example of such an approach is Ambit⁴.

RRAM and Data Access in RRAM

RRAM is a non-volatile memory technology that can store a logical value in terms of its resistance. Typically, an RRAM device is fabricated using a Metal-Insulator-Metal (MIM) structure, where the insulator is rich in oxygen ions. The voltage across the RRAM device controls the drift of the oxygen ions, allowing the formation/rupture of conducting filaments, and thereby forming distinct resistance states. The small diameter of the conducting filaments makes RRAM technology *extremely scalable*, and it is thus an attractive candidate for a future low-cost memory technology. Conventionally, RRAM has two resistance states: the low resistance state (LRS), which represents logical 1, and the high resistance state (HRS), which represents logical 0.

The write operation in RRAM is performed in two phases by applying write voltages via wordlines (WLs) and bitlines (BLs) – SET (to write 1's) and RESET (to write 0's). Read operations in RRAM are performed by applying a voltage pulse of typically lower magnitude than the write voltage and measuring the current flowing out of the cell using current-based sense amplifiers. In addition to performing memory operations, RRAM can also execute logical functions by realizing unique circuit connections and applying voltages. In this article, we focus on Memristor Aided loGIC (MAGIC)² to perform computation using RRAM cells. Advantages of MAGIC include crossbar compatibility, statefulness (which enables true PIM^{6,11}), functional completeness, and parallel execution⁶. Bitwise SIMD MAGIC NOR and NOT operations can be performed inside RRAM to realize PIM merely by applying a voltage pattern, and without any modification of the conventional memory crossbar structure.

Why a New Memory Controller for RRAM?

RRAM and DRAM differ in various ways that affect the memory controller:

- **Row-oriented vs. column-oriented access mechanism.** The DRAM access mechanism is row-oriented, where an access starts with opening a memory row and reading its data to the row buffer. Note that this operation is performed irrespective of whether the request type is read or write. In contrast, RRAM uses different system resources for different operations (*i.e.*, voltage drivers during write vs. voltage drivers + sense amplifiers during read) as well as unique biasing schemes (*i.e.*, half-select during write vs. floating during read). This means that *the type of access* must be known in advance (which is sent using the column command). Hence, the RRAM access protocol needs to be column-oriented (meaning centered around read/write (column) command rather than a row command) as opposed to row-oriented in DRAM. Note that while we propose to change what RRAM protocol should be centered around, each row-access (*i.e.*, opening a memory row and reading data to the row buffer) and column-access (*i.e.*, reading/writing to few columns) have the same meaning as DDRx.
- **Open-page vs. closed page policies.** RRAM multiplexes multiple BLs to a single sense amplifier to reduce the cost of the chip⁵. As a result, unlike DRAM, data from a full row is not available at the row buffer during a read operation. Consequently, the RRAM




Figure 1. Execution time of RRAM accessed using the DDR4-2400 based memory controller normalized to DRAM accessed using a similar memory controller.

memory controller cannot employ an open-page policy and generally, a restricted closed-page policy is used (*i.e.*, closing the row immediately after an access is serviced).

- **Refresh and restore.** A few DRAM operations, namely restore and refresh, are not required in RRAM because of non-destructive reads and high retention times of RRAM.
- **PIM support.** Since DRAM cells cannot perform PIM, the traditional memory controller does not support PIM instructions. Thus, to exploit PIM capabilities of RRAM, we need to support PIM instructions for the RRAM memory controller.

To demonstrate the disadvantages of accessing RRAM using the DRAM protocol, we model an RRAM-based main memory as described in the methodology section, and access it using DDR4-2400. Fig. 1 shows execution time for different workloads listed in Table 1. For these experiments, we use a restricted closed-page policy and the same auxiliary timing constraints as DDR4 (*i.e.*, tCCD, tRRD, tFAW, and tWRT)⁷. We model RRAM read and write latencies by tCAS and tWR, respectively. RRAM accessed using a DDRx-based memory controller is 66.3% slower than DRAM. Additionally, we will show in the next sections that the access protocol also affects these values. Finally, we will show how customizing the memory controller for RRAM can unleash its full potential.

CONCEPT DESIGN

In this section, we propose a memory controller to perform efficient memory as well as PIM operations in RRAM. The controller is divided into the front-end (responsible for receiving and storing the memory requests and PIM commands such as MAGIC NOR/NOT) and the back-end (responsible for converting requests into low-level technology-specific commands and sending these commands to access main memory data while maintaining timing constraints). We present micro-architectural modifications of transaction queues to support new PIM instructions. Then, we discuss our proposed RRAM access protocol and the scheduling algorithm.

CONCEPT Front-End Design

On a last-level cache miss, a request is transferred to the memory controller to perform a memory transaction, and the front-end of the memory controller is responsible for receiving this request. First, the physical address is decoded to a memory-specific coordinate in terms of channel, rank, bank, row, and column. In this work, we assume a traditional scheme of address decoding where upper bits of the physical address are mapped to rows to exploit row-buffer locality, lower bits are mapped to columns, and middle bits are mapped to bank addresses. We then propose changes to modify the transaction queue structure.

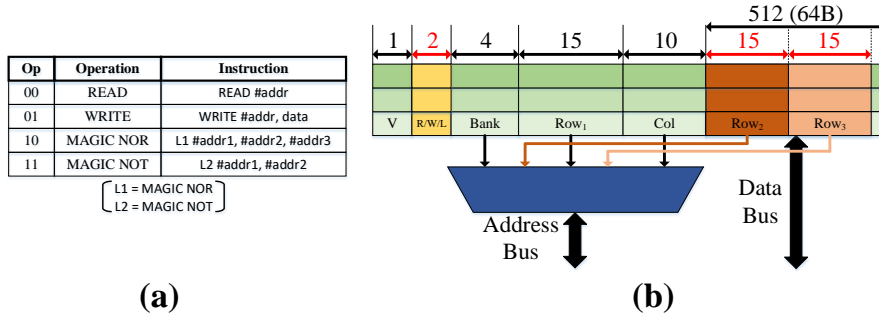


Figure 2. (a) ISA extension of the memory controller to support PIM instructions using MAGIC, and (b) modified design of the transaction queue that can also support PIM instructions. The different field lengths correspond to 4 Gb x8 DDR4 DRAM. Legends: V-valid, Op-opcode. Bit widths in red font represent the proposed extension from a traditional transaction queue design.

Fig. 2a shows the extended ISA of our memory controller. The first two entries in the table are conventional read and write instructions – a read instruction with an address field and a write instruction with an address as well as a data field. The last two entries are the MAGIC instructions to perform SIMD NOR and NOT operations. These instructions are unique in that they use higher number of addresses than read/write without data. MAGIC NOR instruction uses three addresses - two for inputs and one for output. MAGIC NOT instructions use two addresses only – one for inputs and one for output. To incorporate PIM instructions in transaction queues, the fields of the transaction queue can be naively extended to support three addresses independently. However, this solution would incur significant hardware overhead, and thus is impractical. We propose to incorporate the PIM instruction format in transaction queue structure with minimum area overhead. Our proposal is based on four key observations.

- To reduce the pin-count (and hence, the packaging cost), the memory controller sends row and column addresses in time-multiplexed fashion.
- The data field is much wider (*i.e.*, 512-bits) than the address field (*i.e.*, 15-bits for a row address).
- MAGIC operations can be performed only on vectors present in the same RRAM bank (at row granularity), so only one bank address is required per instruction.
- PIM instructions consist only of addresses and not data since they are performed on operands already situated inside the memory.

We thus propose to reuse a few of the data-field bits to create a complete MAGIC instruction within the queue without significantly increasing storage space. Figure 2b shows the design of a modified transaction queue where two row addresses are replicated within the unused data field and a path to the multiplexer is created for time-multiplexing them. Here, we assume that the PIM instructions are executed row-wise; hence, three (two) row addresses and one bank are required to execute a MAGIC NOR (NOT) instruction on a bit-vector present across a whole row. Although PIM operations can be performed at finer granularity by selecting/deselecting different columns/rows for execution³, this approach complicates the control logic and peripheral circuit, and hence, in this article, we focus on exploiting full parallelism from a row, assuming massive parallelism requirement from the workload. Note that even if PIM accesses are performed using more row address bits, it is always possible to accommodate two row addresses in the wide data field. This will minimize the hardware overhead of incorporating PIM instructions with more than one address. The only insignificant overhead is because of a somewhat complex instruction decoder, an additional bit for two-bit opcode (Fig. 2b), and decoding logic for the multiplexer.

CONCEPT Back-End Design

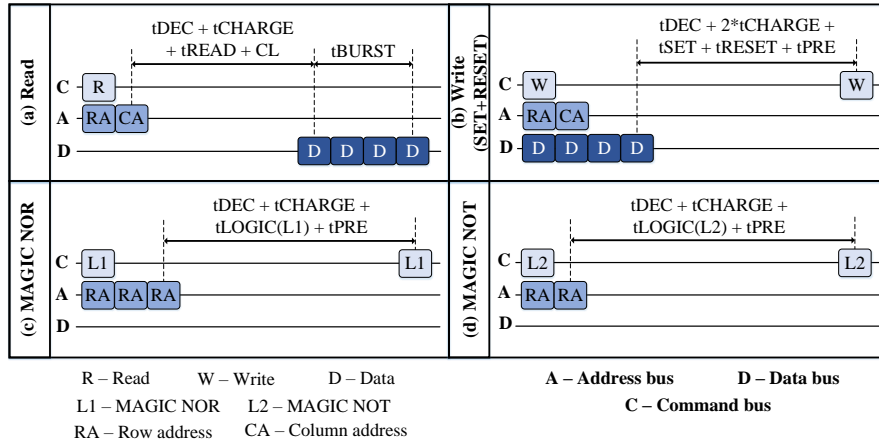


Figure 3. Timings and bus activities of different operations in R-DDR -- (a) read operation, (b) write operation (composed of SET and RESET phases), (c) MAGIC NOR (L1) operation, and (d) MAGIC NOT (L2) operation.

Once the memory requests are latched at the transaction queue, the back-end of the memory controller is responsible for splitting the requests into technology-specific commands and sending these commands at specific time intervals, in accordance with the timing constraints. In this subsection, we present the R-DDR protocol: an RRAM access protocol similar to DDRx, which can incorporate both memory and PIM instructions.

R-DDR: An RRAM Access Protocol Similar to DDR

Since read, write, and logic operations utilize unique resources on the memory die, we propose a column-oriented protocol R-DDR to access RRAM instead of a row-based protocol. To service a read request from the queue using this protocol, the memory controller sends a read command first rather than sending a generic command like *activate*. This prepares the memory to initiate a memory request since different hardware blocks need to be invoked for different requests. Furthermore, since both row and column addresses are required to be sent, R-DDR sends both the addresses in time multiplexed fashion one after the other to avoid the pin-count increase. This is necessary to bias the correct WLs/BLs for write and logic operations, as well as to multiplex the correct BLs to the sense amplifiers. Once the row and column addresses have been sent along with the type of access, the memory starts servicing the access.

When no memory request is being serviced, we say that R-DDR is in the idle state. To eliminate leakage power in the idle state, we choose to bias all the WLs/BLs to a ground voltage so that no current flows through the arrays and data does not get corrupted because of voltage jitters.

Read operations in R-DDR are performed by activating the target WL with a read voltage and multiplexing target BLs to be read to sense amplifiers by means of a column address. Hence, current flows from the driver and RRAM cells to the sense amplifiers, where it is converted to a full logical value. Figure 3a shows the timing diagram of a read operation where there are four latencies associated reading data out from the memory cells and transferring it to the I/O pads once all the addresses are received: (1) address decoding latency (*i.e.*, t_{DEC}); (2) RC-delay of wires to bias WL to target voltage (*i.e.*, t_{CHARGE}); (3) sampling and reading current flowing through RRAM via sense amplifiers (*i.e.*, t_{READ}); and (4) the latency to transfer data from the sense amplifiers to the I/O pads (*i.e.*, CL/t_{CAS} , which is like DRAM). After a time delay equal to the sum of these four latencies, we can expect a data burst to start at the data bus, after which, it takes four clock cycles (*i.e.*, t_{BURST}) for all 64B to get transferred to the CPU.

Write operations in R-DDR are performed by activating target WLs and BLs to write voltages in two steps⁷ - SET and RESET. The R-DDR protocol initiates a write request by sending the write command along with row and column addresses and the data that needs to be written. The timing diagram of write operations is shown in Figure 3b. There are five latencies associated with write operations: (1) address decoding latency (*i.e.*, tDEC); (2) the RC-delay to bias selected WLs and BLs for SET process (*i.e.*, tCHARGE); (3) the delay to SET the device itself (*i.e.*, tSET) and charge the array again for the RESET operation (*i.e.*, tCHARGE - required since SET and RESET operations are performed using different voltage levels and polarity and they are applied to different cells), (4) the time to RESET the device (*i.e.*, tRESET); and finally, (5) the delay to precharge the array to the idle state to prepare for next access (*i.e.*, tPRE).

Logic operations in R-DDR are performed like write operations. As shown in Figure 3(c,d), a MAGIC NOR (NOT) operation is initiated by sending a command and three (two) addresses for inputs and outputs. The delay associated with this operation is comprised of the delay to decode addresses (*i.e.*, tDEC), the RC-delay to charge wires to appropriate values (*i.e.*, tCHARGE), time to perform logic itself (*i.e.*, tLOGIC), and the delay to precharge the array to the idle state to prepare for the next operation.

Consecutive requests of the same type to different rows in a bank are handled by first closing and precharging the row of the preceding request to service the next one. A more interesting case occurs when two consecutive requests access the same row. In DRAM, both reads, and writes are performed using a row buffer that is generally much wider than a single access. Because of this over-provisioning, once the data is read at the row buffer (which incurs high latency of tRCD), subsequent requests to the same row can be serviced with a much lower latency (*i.e.*, tCCD). Conversely, in RRAM, row buffer is only used in read operations. In case of back-to-back write and logic requests to the same bank, even if they access the same row/column, the array needs to be precharged to service the next request. Moreover, although row buffers are used in reads, not a full row is read out because of bulky sense amplifiers, and hence, most RRAM designs immediately precharge a row to service next read. However, there is still some over-provisioning while reading data out. We propose to exploit this limited locality to service multiple read requests to the same row with lower latency when they hit on the narrow row buffer.

Consecutive requests of different types: when a read request follows a write request, DRAM specifies a timing constraint called tWTR, *i.e.*, a write-to-read bus turn-around delay. Although separate I/O paths exist for read and write requests, tWTR is non-zero because of internal bus and row buffer sharing between read and write requests. This gives rise to a long write-to-read turn-around time (*i.e.*, 7.5 ns), which reduces DRAM bandwidth. There are two ways to address this problem: by having independent GIO (global I/O) buses for reads and writes that can reduce tWTR, or by buffering write requests to clear the way for reads. The latest DRAM architectures (such as HBM2 - High Bandwidth Memory) use independent I/O lines - RGIO and WGIO – to increase core frequency, and we also expect RRAM vendors to adopt such a design to alleviate bottlenecks due to tWTR. Nevertheless, even if this is not the case, internal bus contention can be avoided by buffering write requests on the memory chip to delay the write request and allow performance-sensitive read requests to get serviced.

To exploit internal parallelism, DRAM can access multiple banks in parallel concurrently. Although such parallelism is desirable, it increases the amount of current required to service many requests, increasing the power consumed in the chip. Furthermore, since the memory market is extremely cost-sensitive, DRAM is fabricated using a process with only three metal layers, which can only withstand a certain amount of power. Therefore, in practice, parallelism in DRAM is regulated by two timing parameters – tRRD (*i.e.*, row-to-row delay) and tFAW (*i.e.*, four-bank activation window). To control instantaneous peak power at any time, activate commands to different banks are separated by tRRD. Furthermore, tFAW is defined to specify a rolling time-frame in which a maximum of four banks can be activated concurrently. This is to control the current drawn from the source in a time window of tFAW, and generally, $tFAW > 4 * tRRD$. Values of these parameters change according to the geometry of the DRAM device, which changes the activation current. RRAM can be fabricated in metal layers on a standard logic process and trades off cost with performance by not requiring transistors as access elements. We expect the logic layer in

Table 1. SPEC CPU 2006 workload setup.

#	MPKI	Benchmark Mixture
M1	Low- Low- Mid- High	povray-gromacs-soplex-cactusADM
M2		GemsFDTD-gamess-wrf-milc
M3		hmmmer-h264ref-perlbench-mcf
M4		sphinx3-namd-gcc-lbm
M5		tonto-dealII-bwaves-leslie3d
M6	Low- Mid- Mid- High	calculix-libquantun-zeusmp-cactusADM
M7		xalancbmk-soplex-wrf-milc
M8		gobmk-wrf-perlbench-mcf
M9		sjeng-gcc-bwaves-lbm
M10		omnetpp-libquantum-zeusmp-leslie3d
M11	Low- Mid- High- High	bzip2-soplex-cactusADM-milc
M12		astar-wrf-mcf-lbm
M13		povray-perlbench-leslie3d-cactusADM
M14		gromacs-gcc-milc-mcf
M15		GemsFDTD-bwaves-lbm-leslie3d

RRAM to support much higher power as compared to a cheap DRAM process, and as a result, the tFAW to be constraint relaxed (*i.e.*, $tRRD=4*tFAW$) in RRAM.

Although implementing these optimizations in RRAM incurs an increased cost, this cost is marginal considering the benefits that would be obtained by achieving DRAM-like performance in RRAM. Furthermore, the cost-per-bit for RRAM is lower than for DRAM. While it is also possible to perform power and parallelism-related optimizations to improve its performance in DRAM as mentioned above, we believe there is no justifiable reason for doing so, given the increased cost, and given that DRAM performs well without these optimizations.

EXPERIMENTAL METHODOLOGY

To demonstrate the benefits of our memory design, we compare the performance and energy consumption of our memory system with (a) DDR4-accessed RRAM and (b) DDR4-2400 DRAM. Furthermore, to present the benefit of supporting MAGIC execution from the memory controller,

Table 2. Parameter values used for simulation.

Parameter	Value	Parameter	Value
V_READ	0.7V	tREAD	11nCK
V_SET	1.4V	tSET	27nCK
V_RESET	1.4V	tRESET	27nCK
V_MAGIC_NOR	1.8V	tMAGIC_NOR	35nCK
V_MAGIC_NOT	2V	tMAGIC_NOT	35nCK
E_READ (array)	0.84pJ/bit	E_READ (periphery)	4.8pJ/bit
E_SET	6.9pJ/bit	E_SET (periphery)	4.8pJ/bit
E_RESET	6.9pJ/bit	E_RESET (periphery)	4.8pJ/bit
Roff	175kOhm	Ron	7kOhm
tDEC	1nCK	tPRE	1nCK
tCHARGE	1nCK	CL	17nCK
tDEC	1nCK	tBURST	4nCK
tWTR	0nCK	tFAW	16nCK
nCK	0.833ns	tRRD	4nCK

we compare our PIM design with (a) DDR4-2400 based CPU model, and other DDRx-based similar PIM designs, *i.e.*, (b) Pinatubo⁵, and (c) Ambit⁴.

We choose a HfOx-based device⁸ for our comparison of RRAM and use VTEAM device model¹³ to fit its parameters to the experimental data. Using the voltage levels listed in Table 2, we estimate the array-level latency and energy of read, write and MAGIC operations using SPICE. Furthermore, we use NVSim¹⁴ to determine the latency and energy of peripheral circuit (which also includes transferring data to and from sense amplifiers and I/O pads).

To evaluate memory operations, we use NVMain¹⁵. We use SniperSim to gather the main memory traces for the SPEC CPU2006 benchmark suite for 500M instructions, after 100M instructions to warm up the cache. The CPU consists of a 4-core, 2.4 GHz Nehalem processor with 2 MB/core, 16-way LRU LLC. DRAM standard is DDR4-2400 17-17-17¹⁰, with a capacity of 4 Gb with a single channel, a single rank, 16-banks, and 8KB page size. The workload setup used to evaluate memory system is listed in Table 1. Note that nCK represented the numbers of clock cycles, and the clock period is 0.833ns (according to DDR4-2400 I/O speed).

To evaluate PIM operations, we use bitmap indices to track users' characteristics⁴. The database stores the gender and the per-day log-in activity of millions of users in memory. We run two queries: (a) how many unique users were active every week for the past w weeks, and (b) how many male users were active each of the past w weeks. To execute these queries, $6w$ bitwise OR, $(2w-1)$ bitwise AND and $(w+1)$ bitcount operations need to be performed for all u -users. For all PIM approaches, we accelerate bitwise AND/OR operations using PIM, and bitcount operations are performed on the CPU. We build an in-house simulator to compare performance of the different approaches, where we model latency of PIM operations and data transfer between the CPU and main memory. For all PIM approached, the same I/O frequency and memory organization are used as the baseline (*i.e.*, 4Gb DDR4-2400). Ambit and Pinatubo numbers are adopted from the original articles and are extended for the configuration used for comparison. Since our workload is data bound, we ignore the time required to compute data on the CPU.

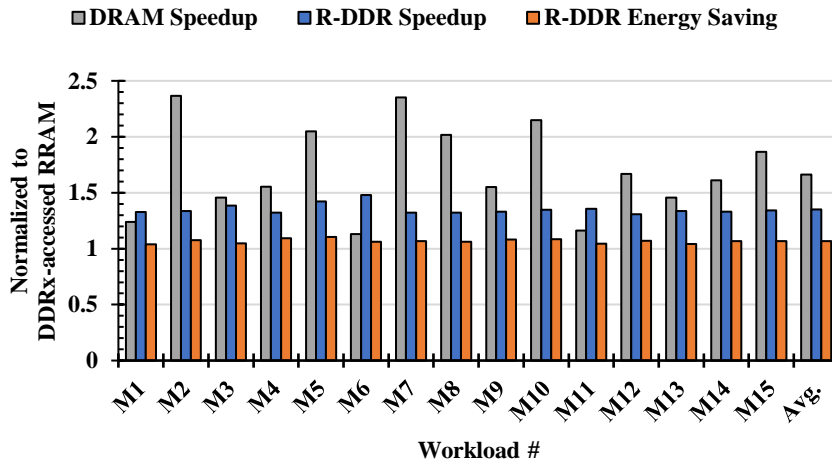


Figure 4. IPC speedup and energy savings of R-DDR normalized to performance and energy consumption of RRAM that is accessed using DDR4-2400 (higher is better). Write intensity (*i.e.*, fraction of write requests in all requests) is also shown for all workloads.

RESULTS

Memory Performance and Energy

Fig. 4 compares performance and energy efficiency of three approaches – (a) DDRx-accessed RRAM (baseline), (b) DRAM, and (c) R-DDR accessed RRAM (this work). R-DDR gains 35% speedup and 7% energy savings on average compared to baseline. The reasons for the improved performance include exploitation of partial row-buffer locality, removal of unnecessary steps such as refresh and restore, and redesign of timing constraints similar to DDRx, but specifically for RRAM. The exploitation of partial row-buffer locality is a major contributor to the superior performance, as it allows us to perform read operations at a much lower latency. In addition, relaxing auxiliary timing constraints such as tWTR and tFAW contributes to the performance.

Energy benefits are thanks to the ability to exploit partial row-buffer locality and reduced execution time. Employing a narrow row-buffer policy reduces the number of reads that must be performed from the RRAM array, hence saving the read energy of requests hitting on the row buffer. Furthermore, reduction in execution time of workload results in lower background energy. However, energy savings of R-DDR is lower than performance improvement compared to baseline because energy consumption is dominated by write requests.

PIM Performance

Fig. 5a shows the normalized speedups of Pinatubo, Ambit, and the proposed MAGIC-based PIM approach compared to a von Neumann machine for different numbers of weeks and users. For a fair comparison, we include the execution time to compute end-to-end workload and ignore cycles to send commands from the memory controller to invoke PIM operations since the overhead of sending commands that take a couple of cycles is negligible compared to PIM operation that takes several tens-hundreds of cycles, and it can be further minimized by pipelining PIM execution with sending commands. Note that although MAGIC is proposed previously, the benefits of MAGIC are enabled because of system proposed modifications without which, it is not possible to realize MAGIC. The figure shows that Pinatubo performs similar to the CPU-based approach; Ambit and our system outperforms CPU by approximately 5X. Pinatubo’s low performance is attributed to its limited compute parallelism due to sense amplifier multiplexing. Both Ambit and MAGIC perform PIM operations in parallel on vectors of size of a full row with similar compute bandwidth. Moreover, since this workload employs $u(w+1)$ bit-read operations to the CPU for bitcount, data transfer between CPU and memory takes around 90% of the execution time in both architectures.

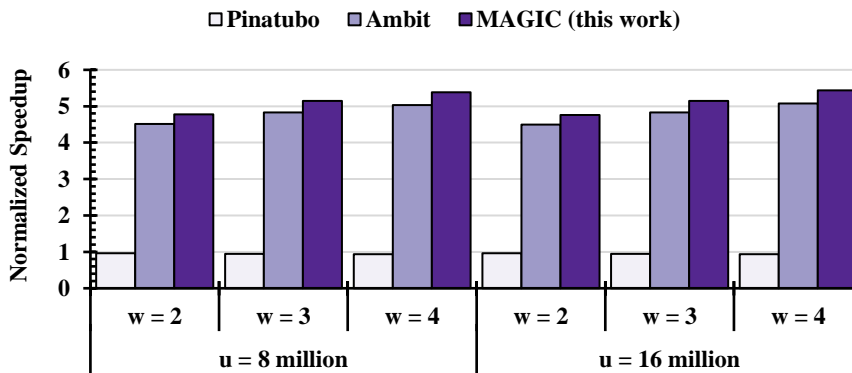


Figure 5. Speedup of Pinatubo, Ambit, and proposed MAGIC-based PIM normalized to von Neumann-based computing machine for query execution in bitmap indices (higher is better). Legends: w - #weeks, u - #users.

Hence, with slightly higher MAGIC latency, MAGIC still performs similar to Ambit. However, benefits of MAGIC approach will be more evident for workloads with greater opportunity to reduce data transfer.

CONCLUSIONS

RRAM has the potential to solve the DRAM scaling problem and the von Neumann bottleneck. However, merely replacing DRAM with RRAM is not the answer. In this article, we propose CONCEPT – a column-oriented memory controller for RRAM to optimize its data access pattern as well as exploit PIM capabilities. We show that a memory controller designed around the technological constraints of RRAM can achieve close-to-DRAM performance. Furthermore, supporting PIM operations in RRAM can result in a 5X speedup compared to a state-of-the-art CPU-memory based von Neumann model.

ACKNOWLEDGEMENTS

This research is partially supported by the European Research Council under the European Union's Horizon 2020 Research and Innovation Programme (grant agreement no. 757259), by the Viterbi Fellowship at the Technion Computer Engineering Center, by the EU ICT COST Action IC1401, and by the Israel Science Foundation grant no. 1514/17.

REFERENCES

1. A. Pedram, S. Richardson, M. Horowitz, S. Galal, and S. Kvatinsky, "Dark Memory and Accelerator-rich System Optimization in the Dark Silicon Era," *IEEE Design Test*, Vol. 34, No. 2, pp. 39–50, April 2017.
2. S. Kvatinsky, D. Belousov, S. Liman, G. Satat, N. Wald, E. G. Friedman, A. Kolodny, and U. C. Weiser, "MAGIC - Memristor-Aided loGIC," *IEEE Transactions on Circuits and Systems II: Express Briefs*, Vol. 61, No. 11, pp. 895–899, September 2014.
3. N. Talati, S. Gupta, P. Mane, and S. Kvatinsky, "Logic Design Within Memristive Memories Using Memristor-Aided loGIC (MAGIC)," *IEEE Transactions on Nanotechnology*, Vol. 15, No. 4, pp. 635–650, July 2016.
4. V. Seshadri, et al., "Ambit: In-memory Accelerator for Bulk Bitwise Operations Using Commodity DRAM Technology," in *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture*, pp. 273–287, October 2017.
5. S. Li, C. Xu, Q. Zou, J. Zhao, Y. Lu, and Y. Xie, "PINATUBO: A Processing-in-memory Architecture for Bulk Bitwise Operations in Emerging Non-volatile Memories," in *Proceedings of the 53rd Annual Design Automation Conference*, June 2016.
6. J. Reuben, R. Ben-Hur, N. Wald, N. Talati, A. H. Ali, P. E. Gaillardon, and S. Kvatinsky, "Memristive logic: A Framework for Evaluation and Comparison," in *2017 27th International Symposium on Power and Timing Modeling, Optimization and Simulation (PATMOS)*, pp. 1–8, September 2017.
7. X. Cong, D. Niu, N. Muralimanohar, R. Balasubramonian, T. Zhang, S. Yu, and Y. Xie, "Overcoming the Challenges of Crossbar Resistive Memory Architectures," in *Proceedings of the IEEE International Symposium on High Performance Computer Architecture*, pp. 476–488, February 2015.
8. H. Y. Lee, et al., "Evidence and Solution of Over-RESET Problem for HfOx Based Resistive Memory with Sub-ns Switching Speed and High Endurance," in *2010 International Electron Devices Meeting*, pp. 19.7.1–19.7.4, Dec 2010.
9. J. Zhou, K. H. Kim, and W. Lu, "Crossbar RRAM Arrays: Selector Device Requirements During Read Operation," *IEEE Transactions on Electron Devices*, Vol. 61, No. 5, pp. 1369–1376, May 2014.
10. Micron, "4Gb one-die DDR4 SDRAM Datasheet Rev. J," July 2017.

11. A. Haj Ali, R. Ben Hur, N. Wald, R. Ronen, and S. Kvatinsky, "Not in Name Alone: A Memristive Memory Processing Unit for Real In-Memory Processing", *IEEE Micro*, Vol. 38, No. 5, pp. 13-21, September/October 2018.
12. S. Kvatinsky, M. Ramadan, E. G. Friedman, and A. Kolodny, "VTEAM – A General Model for Voltage Controlled Memristor", *Transactions on Circuits and Systems II: Express Briefs*, Vol. 62, No. 8, pp. 786-790, August 2015.
13. S. Kvatinsky, M. Ramadan, E. G. Friedman, and A. Kolodny, "VTEAM – A General Model for Voltage Controlled Memristor", *Transactions on Circuits and Systems II: Express Briefs*, Vol. 62, No. 8, pp. 786-790, August 2015.
14. X. Dong, C. Xu, Y. Xie and N. P. Jouppi, "NVSim: A Circuit-Level Performance, Energy, and Area Model for Emerging Nonvolatile Memory," in *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Vol. 31, No. 7, pp. 994-1007, July 2012.
15. M. Poremba and Y. Xie, "NVMain: An Architectural-Level Main Memory Simulator for Emerging Non-volatile Memories," 2012 *IEEE Computer Society Annual Symposium on VLSI*, Amherst, MA, pp. 392-397, September 2012.
16. E. Carlson, W. Heirman, S. Eyerman, I. Hur, and L. Eeckhout, "An Evaluation of High-level Mechanistic Core Models," *ACM Transactions on Architecture and Code Optimization (TACO)*, October 2014.

ABOUT THE AUTHORS

Nishil Talati is a PhD student at the Computer Science & Engineering Department at the University of Michigan, Ann Arbor. He received an MSc degree in Electrical Engineering from the Technion – Israel Institute of Technology in 2018, and a B.Eng. degree in Electrical Engineering from BITS Pilani, India in 2016. His current research interests include computer architecture, main memory systems, and emerging memory technologies. Contact him at talatin@umich.edu.

Heonjae Ha is a Ph.D. student in Electrical Engineering at Stanford University. He received his B.Eng. degree in Electronics Engineering from Korea University in 2006 and M.S. degree in Electrical Engineering from Stanford University in 2009. Prior to joining the Ph.D. program, Heonjae worked in SK Hynix from 2009 to 2013 where his main role was to design and validate mobile DRAMs. His current research is focused on understanding and improving the energy efficiency of DRAMs. Contact him at hunjaeha@stanford.edu.

Ben Perach received the B.Sc. degree in Mathematics from The Hebrew University of Jerusalem, Jerusalem in 2010 and the M.Sc. degree in Electrical Engineering from Tel Aviv University in 2017. He is currently working toward the Ph.D. degree in Electrical Engineering at the Technion, Israel Institute of Technology. His current research interests include computer architecture with a focus on processor design, and field-programmable gate arrays, security, and data networks. Contact him at bperach@gmail.com.

Ronny Ronen is a Senior Researcher at the Andrew and Erna Viterbi Faculty of Electrical Engineering, Technion – Israel Institute of Technology. He received the B.Sc. and M.Sc. degrees in Computer Science in 1978 and 1979, respectively, both from the Technion – Israel Institute of Technology. Ronny was at Intel from 1980 to 2017 in various technical and managerial positions. In his last role, Ronny led the Intel Collaborative Research Institute for Computational Intelligence (ICRI-CI). Until 2011, Ronny was a senior staff computer architect in the Intel Development Center in Haifa and before that the director of Microarchitecture Research in that center. In these roles, Ronny led/was involved in the initial definition and pathfinding of major leading-edge Intel processors. Earlier, Ronny led the development of several system software products and tools including the Intel Pentium processor performance simulator and several compiler efforts. Ronny holds over 70 issued patents and has published over 20 papers. Ronny is an IEEE Fellow and was an Intel Senior Principal Engineer. Contact him at ronny.ronen@technion.ac.il.

Shahar Kvatinsky is an assistant professor at the Andrew and Erna Viterbi Faculty of Electrical Engineering, Technion – Israel Institute of Technology. He received the B.Sc. degree in computer engineering and applied physics and an MBA degree in 2009 and 2010, respectively, both from the Hebrew University of Jerusalem, and the Ph.D. degree in electrical engineering from the Technion – Israel Institute of Technology in 2014. From 2006 to 2009 he was with Intel as a circuit designer and was a post-doctoral research fellow at Stanford University from 2014 to 2015. Kvatinsky is an editor of Microelectronics Journal and has been the recipient of the 2015 IEEE Guillemin-Cauer Best Paper Award, 2015 Best Paper of Computer Architecture Letters, Viterbi Fellowship, Jacobs Fellowship, ERC starting grant, the 2017 Pazy Memorial Award, the 2014 and 2017 Hershel Rich Technion Innovation Awards, 2013 Sanford Kaplan Prize for Creative Management in High Tech, 2010 Benin prize, and six Technion excellence teaching awards. His current research is focused on circuits and architectures with emerging memory technologies and design of energy efficient architectures. Contact him at shahar@ee.technion.ac.il.