

Family Name:

First Name:

Student ID:

Final Exam

Computer Architecture (227-2210-00L)

ETH Zürich, Fall 2021

Prof. Onur Mutlu

Problem 1 (90 Points):	DRAM Latency	
Problem 2 (85 Points):	RowHammer	
Problem 3 (120 Points):	Processing-near-Memory	
Problem 4 (105 Points):	Memory Access Patterns	
Problem 5 (105 Points):	Emerging Memory Technologies	
Problem 6 (80 Points):	Asymmetric Multicore	
Problem 7 (115 Points):	Interconnects	
Problem 8 (BONUS: 70 Points):	Prefetching	
<hr/>		
Total (770 (700 + 70 bonus) Points):		

Examination Rules:

1. Written exam, 180 minutes in total.
2. No books, no calculators, no computers or communication devices. 10 single-sided (or 5 double-sided) A4 pages of handwritten notes are allowed.
3. Write all your answers on this document, space is reserved for your answers after each question. Blank pages are available at the end of the exam. Do not answer questions on them.
4. Clearly indicate your final answer for each problem. Answers will only be evaluated if they are readable.
5. Put your Student ID card visible on the desk during the exam.
6. If you feel disturbed, immediately call an assistant.
7. Write with a black or blue pen (no pencil, no green, red or any other color).
8. Show all your work. For some questions, you may get partial credit even if the end result is wrong due to a calculation mistake. If you make assumptions, state your assumptions clearly and precisely.
9. Please write your initials at the top of every page.

Tips:

- **Be cognizant of time. Do not spend too much time on one question.**
- **Be concise. You may be penalized for verbosity.**
- **Show work when needed. You will receive partial credit at the instructors' discretion.**
- **Write legibly. Show your final answer.**

Initials: _____

Computer Architecture

December 17th, 2021

This page intentionally left blank

1 DRAM Latency [90 points]

You would like to understand the configuration of the DRAM subsystem of a computer using reverse engineering techniques. Your current knowledge of the particular DRAM subsystem is limited to the following information:

- The physical memory address is 24 bits.
- The DRAM subsystem consists of a single channel, 8 banks, and 1024 rows per bank.
- The DRAM is byte-addressable.
- The most-significant 3 bits of the physical memory address determine the bank. The following 10 bits of the physical address determine the row.
- The DRAM command bus operates at 500 MHz frequency.
- The memory controller uses the open-row DRAM row management policy.
- The memory controller uses first-ready first-come first-serve (FRFCFS) request scheduling policy. With this policy, in every clock cycle, the memory controller picks the oldest request (in the request buffer) for which the corresponding bank is ready to receive the DRAM command needed to serve the request. For example, if the oldest request needs to precharge bank X but the bank is not yet ready to be precharged (i.e., ACTIVATE-to-PRECHARGE latency is not yet satisfied), the memory controller checks the next oldest request in the request buffer. If none of the DRAM banks corresponding to the requests in the request buffer are ready, the memory controller does not issue any command in that cycle.
- The memory controller services requests in order with respect to each bank. In other words, for a given bank, the memory controller first services the oldest request in the *request buffer* that targets the same bank.

You realize that you can observe the memory requests that are waiting to be serviced in the request buffer. At a particular point in time, you take the snapshot of the request buffer and you observe the following requests in the request buffer:

decreasing age ↓	Read 0x660000 ← oldest
	Read 0x460000
	Read 0x640000
	Read 0xA60000 ← youngest

At the same time you take the snapshot of the request buffer, you start probing the DRAM command bus. You observe the DRAM command type and the cycle (relative to the first command) in which the command is seen on the DRAM command bus. The following are the DRAM commands you observe on the DRAM bus while the requests above are serviced. Note that, although the memory controller uses the open-row policy, *a bank might be in precharged state (i.e., closed)* prior to the execution of the command sequence below if no rows are accessed in the bank after the system is powered up. Assume that all banks are ready for a DRAM command (i.e., no bank processes a previously issued command) at cycle 0.

```

Cycle 0 --- PRECHARGE
Cycle 1 --- ACTIVATE
Cycle 2 --- ACTIVATE
Cycle 4 --- ACTIVATE
Cycle 9 --- READ
Cycle 10 --- READ
Cycle 12 --- READ
Cycle 30 --- PRECHARGE
Cycle 34 --- ACTIVATE
Cycle 42 --- READ

```

Answer the following questions using the information provided above.

- (a) [10 points] What is the size (in KB) of a single row? Explain.

- (b) [30 points] What are the following DRAM timing parameters used by the memory controller, in terms of nanoseconds? If there is not enough information to infer the value of a timing parameter, write *unknown* and explain why an exact value cannot be determined. *Assume that READ-to-PRECHARGE latency is 1 cycle.*

- i) ACTIVATE-to-READ latency

- ii) ACTIVATE-to-PRECHARGE latency

- iii) PRECHARGE-to-ACTIVATE latency

- (c) [20 points] What is the row buffer state of the banks *prior* to the execution of any of the above requests? Fill in the table below indicating whether a bank has an open row or not. Write *unknown* in case the provided information is insufficient to determine the row buffer state.

	Open or Closed?
Bank 0	
Bank 1	
Bank 2	
Bank 3	
Bank 4	
Bank 5	
Bank 6	
Bank 7	

- (d) [30 points] To improve performance, you decide to implement the idea of *Tiered-Latency DRAM (TL-DRAM)* in the DRAM chip. Assume that a bank consists of a single subarray. With TL-DRAM, an entire bank is divided into a near segment and far segment. When accessing a row in the near segment, the ACTIVATE-to-READ latency *reduces* by 2 cycles and the ACTIVATE-to-PRECHARGE latency reduces by 5 cycles. When precharging a row in the near segment, the PRECHARGE-to-ACTIVATE latency reduces by 1 cycle. When accessing a row in the far segment, the ACTIVATE-to-READ latency *increases* by 1 cycle and the ACTIVATE-to-PRECHARGE latency increases by 2 cycles. When precharging a row in the far segment, the PRECHARGE-to-ACTIVATE latency increases by 2 cycles. The following table summarizes the changes in the affected latency parameters.

Timing Parameter	Near Segment Latency	Far Segment Latency
ACTIVATE-to-READ	-2	+1
ACTIVATE-to-PRECHARGE	-5	+2
PRECHARGE-to-ACTIVATE	-1	+2

Assume that the rows in the near segment have smaller row ids compared to the rows in the far segment. In other words, physical memory row addresses 0 through $N - 1$ are the near-segment rows, and physical memory row addresses N through 1023 are the far-segment rows.

If the above DRAM commands are issued 6 cycles faster with TL-DRAM compared to the baseline (the last command is issued in cycle 36), how many rows are in the near segment, i.e., what is N ? Show your work.



2 RowHammer [85 points]

2.1 RowHammer Attacks

In order to characterize the vulnerability of your DRAM device to RowHammer attacks, you must be able to induce RowHammer bit flips. Assume the following about the target system:

- The CPU has a single-core in-order processor, and does not implement virtual memory.
- The physical memory address is 16 bits.
- The DRAM subsystem consists of two channels, four banks per channel, and 64 rows per bank.
- The memory controller employs the open-row policy.
- The memory controller does not employ any remapping or scrambling schemes for the physical address.
- There is no RowHammer mitigation mechanism implemented in any part of the system.
- All the cells in the DRAM subsystem are equally vulnerable to RowHammer-induced bit flips.

You implement several programs using the instructions shown in Table 1.

Instruction	Description	Functionality
B LABEL	Unconditional Branch	PC = LABEL
STORE IMM, Rs	Store word to memory	MEM[IMM] = Rs
CLFLUSH IMM	Cache line flush	Flush cache line containing IMM

Table 1: Instruction Descriptions.

- (a) [20 points] You run Program 1 below, but you cannot observe any bit flips in the target system. You figured out that the number of activations is much lower than your expectation. Give **two** potential reasons that might cause Program 1 to fail to introduce a sufficient number of row activations to induce RowHammer bit flips.

Program 1

```

1:  LOOP:
2:    STORE 0x8732, R0
3:    STORE 0x8734, R0
4:    CLFLUSH 0x8732
5:    CLFLUSH 0x8734
6:    B LOOP

```

- (b) [30 points] You try Programs 2a, 2b, and 2c, but find that *only two of them can induce RowHammer bit flips* in your DRAM subsystem. Which program segment is the one that does *not* induce

RowHammer bit flips? Justify your answer.

Program 2a

```
1: LOOP:
2:   STORE 0x8732, R0
3:   STORE 0x98CD, R1
4:   CLFLUSH 0x8732
5:   CLFLUSH 0x98CD
6:   B LOOP
```

Program 2b

```
1: LOOP:
2:   STORE 0xF1AB, R0
3:   STORE 0x0054, R1
4:   CLFLUSH 0xF1AB
5:   CLFLUSH 0x0054
6:   B LOOP
```

Program 2c

```
1: LOOP:
2:   STORE 0x2B97, R0
3:   STORE 0xDA68, R1
4:   CLFLUSH 0x2B97
5:   CLFLUSH 0xDA68
9:   B LOOP
```


2.2 RowHammer Defenses

You are given a system that crashes frequently. After a thorough investigation, you notice that this crash happens when a DRAM row (say Row A) is activated many times. The evidence suggests that this failure is caused by RowHammer bit flips. Therefore, you analyze and collect the following information about the RowHammer vulnerability of the DRAM chip, the RowHammer defense mechanism, and the memory access pattern.

RowHammer vulnerability of the DRAM chip. A row needs to be activated at least 20000 times to induce a bit flip in its physically adjacent rows. The *blast radius* is one, meaning that hammering a DRAM row induces bit flips only in its physically adjacent rows.

RowHammer defense mechanism. The RowHammer defense mechanism is implemented in the memory controller. The defense mechanism deterministically and precisely counts the activation count of each row. When the activation count of a Row X reaches 4096, the defense mechanism sends row activation commands to row addresses X-1 and X+1.

Memory access pattern. The access pattern that causes RowHammer bit flips activates Row A 30000 times within a 32ms time window. During these activations, the defense mechanism activates rows A-1 and A+1 once every 4096 activations targeting row A.

- (a) [15 points] Even though the RowHammer defense mechanism seems to refresh rows A-1 and A+1, you still observe RowHammer bit flips. Why is that? Explain.

- (b) [10 points] Would implementing Probabilistic Row Activation (PARA), as proposed in the original RowHammer paper from ISCA 2014, in the memory controller instead of the explained defense mechanism solve the problem? Explain your reasoning.

- (c) [10 points] Would implementing BlockHammer, as discussed in Lecture 16c, in the memory controller instead of the explained defense mechanism solve the problem? Explain your reasoning.

3 Processing-near-Memory [120 points]

You want to accelerate the following two pieces of code from Application 1 (App1) and Application 2 (App2).

```
// App1. Registers are 4-byte wide
movi R1, #0x1000          // Store the base address of A in R1
movi R2, #0x8000          // Store the base address of B in R2
movi R3, #0

Loop:
  ld R4, [R1, R3]         // R4 = MEM[R1 + R3]
  ld R5, [R2, R3]         // R5 = MEM[R2 + R3]
  mult R6, R4, #0xF       // R6 = R4 * 0xF
  add R6, R6, R5          // R6 = R6 + R5
  st [R1, R3], R6         // MEM[R1 + R3] = R6

  inc R3                  // R3++
  bne R3, 1000000000, Loop // If R3 != 1000000000, jump to Loop
```

```
// App2. Registers are 4-byte wide
movi R1, #0x1000 // Store the base address of A in R1
movi R2, #0x8000 // Store the base address of B in R2
movi R3, #0
movi R4, #0

Loop:
  ld R5, [R1, R3]         // R5 = MEM[R1 + R3]
  ld R6, [R1, R3, #4]     // R6 = MEM[R1 + R3 + #4]
  sub R5, R5, R6          // R5 = R5 - R6
  mult R5, R5, R5         // R5 = R5 * R5
  add R4, R4, R5          // R4 = R4 + R5
  sqrt R4, R4             // R4 = sqrt(R4)
  st [R2, R3], R4         // MEM[R2 + R3] = R4

  inc R3, #2              // R3=R3+2
  bne R3, 1000000000, Loop // If R3 != 1000000000, jump to Loop
```

We make the following assumptions about the baseline CPU where both applications run:

- The CPU is a single-issue in-order processor and all load/store operations are serialized, i.e., the latency of multiple memory requests cannot be overlapped.
- The clock frequency of the CPU is 1 GHz.
- Each memory operation (i.e., ld, st) takes 100 ns.
- Each simple arithmetic operation (i.e., add, sub, mult, inc) and branch operation (i.e., bne) takes 1 clock cycle to execute.
- A complex arithmetic operation (i.e., sqrt) takes as long as 50 simple arithmetic operations.
- All memory operations (i.e., ld, st) and arithmetic operations (i.e., add, sub, mult, inc, sqrt) operate on 4 bytes of data.

- (a) [30 points] What is the execution time of App1 and App2 when running on the baseline CPU? Show your work. Hint: Do *not* account for the execution time of the initial `movi` instructions, since it is negligible in comparison to the loops.

App1:**App2:**

You have learned in class that Processing-near-Memory (PnM) architectures can accelerate memory-bound workloads, since they provide higher bandwidth and lower latency than conventional processor-centric architectures. You decide to design PnM accelerators for App1 and App2.

The memory device you use is an early-generation 3D-stacked memory with an internal memory bandwidth of only 40 GB/s (i.e., twice the bandwidth of a single-channel 2D DDR4 memory). The 3D-stacked memory allows you to embed compute resources at the base die of the memory cube, called *logic layer*. However, the logic layer imposes several design limitations:

- The area available to build your accelerator in the logic layer of the 3D-stacked memory is **100 mm²**.
- The maximum clock frequency of the accelerator in the logic layer of the 3D-stacked memory is $\frac{1}{10} \times$ that of baseline CPU.
- The accelerator consists of one or more *processing elements*, each of which contains several *functional units*. Table 1 shows the functional units available to build your accelerators.

- A processing element of the accelerator executes the same computation as an entire iteration of the loop (i.e., *all* the instructions in the loop body). The processing element executes an iteration completely before moving to the next iteration. Therefore, if there are N processing elements, N iterations of the loop are executed in parallel.
- The loop iterations are executed on the processing elements as decided by a compiler, which unrolls the loop and preassigns iterations to the processing elements (i.e., the compiler schedules the iterations statically).
- Each functional unit of a processing element of the accelerator can execute one instruction of the loop body at a time. The same functional unit can execute different instructions (e.g., an arithmetic unit is capable of executing `add`, `sub`, `mult`, `inc`) at different times. Two functional units can execute in parallel, if there is no dependence between the operands.

Table 1: Functional units available to build your accelerators.

Functional Unit	Description	Latency (cycles)	Area (mm^2)
Arithmetic Unit	Arithmetic unit capable of executing <code>add</code> , <code>sub</code> , <code>mult</code> , <code>inc</code>	1	1
Load and Store Unit	Load and store unit. It can issue 1 4-byte <code>ld/st</code> operation per cycle	1	1
Branch Unit	Executes conditional branches (<code>bne</code>) with 100% accuracy	1	1
Square Root Unit	Executes square root (<code>sqrt</code>) operations	70	30

You design your accelerator with the maximum possible number of processing elements, in order to be able to run in parallel as many loop iterations as possible. A processing element can have as many functional units as possible, in order to extract from the code as much Instruction Level Parallelism (ILP) as possible.

- (b) [30 points] What is the area of the configuration of your PnM accelerator that provides the highest performance for each application while fitting in the PnM area budget? Show your work.

App1:**App2:**

- (c) [30 points] Are the PnM accelerators obtained in (b) capable of fully utilizing the memory bandwidth that the 3D-stacked memory provides for App1 and for App2? Show your work.


App1:

App2:

- (d) [30 points] What is the speedup of the PnM accelerators compared to the execution of App1 and App2 on the baseline CPU? Show your work.

App1:

App2:



4 Memory Access Patterns & Interference [105 points]

Consider a graph that consists of V vertices and E edges and is represented using a compression format (CF). CF consists of 3 arrays:

- **source**: holds the number of edges (neighbors) of vertex i at position i .
- **dst**: holds the destination (neighbors) vertices' ids for each source vertex.
- **weight**: holds the weight of the edge connecting the source vertices and the destination vertices.

Figure 1 demonstrates an example graph ($V=5$, $E=7$) represented using the CF format. The vertex with ID = 0 has 2 neighbors: destination vertices 1 and 2.




Figure 1: Graph with $V=5$ vertices and $E=7$ edges

Consider App 1, a graph analytics application that (i) visits every vertex of the graph, (ii) traverses its neighbors and (iii) calculates a property of the neighbor vertex.

All the data structures hold 4-byte integer values.

Graph analytics application (App 1)

```

1 old_rank = malloc(sizeof(int)*graph.vertices); // Data Structures of App 1
2 rank = malloc(sizeof(int)*graph.vertices);
3 dst = malloc(sizeof(int)*graph.edges);
4 weight = malloc(sizeof(int)*graph.edges);
5 source = malloc(sizeof(int)*graph.vertices);
6
7 initialize_data_structures();
8
9 int edge = 0;
10 for( i=0; i < graph.vertices; i++ )
11     neighbors = source[i];
12     for( j = 0; j < neighbors; j++ )
13         rank[dst[edge]] += weight[edge] * old_rank[i];
14     edge += 1;

```

Consider a system with the following configuration:

- The system consists of a single-core CPU and a single DRAM channel with 4 banks.
- All memory requests from the CPU are served by the DRAM. There are no caches in the system.
- Memory requests are serialized, i.e., memory requests do *not* overlap and CPU stalls until data arrives.
- The processor requests data from DRAM **only** to access the 5 data structures used in App 1.
- **No memory access reordering** is performed by the processor. For example, during the first iteration (i.e. $i = 0$, $j = 0$, $edge = 0$) the processor accesses memory in the following order: `source[0]`, `weight[0]`, `old_rank[0]`, `dst[0]`, `rank[dst[0]]` (assume a single request for reading and writing to `rank`), `source[1]`, etc.
- All the data structures are aligned at the start of a bank row.

- All the data structures are mapped in consecutive memory locations in each bank. For example, given a 4KB DRAM row size, rank[0]-rank[1023] is stored in Row 0, rank[1024-2047] in Row 1, etc.
- Latency of activating a row: 100 cycles.
- Latency of reading data from an activated row: 16 cycles.
- Latency of writing data to an activated row: 20 cycles.
- Latency of precharging a row: 80 cycles.
- DRAM Bus width: 4 bytes.
- For every request to DRAM, the processor receives 4 bytes.
- Latency of transferring 4 bytes: 4 cycles.
- Row buffer size: 2KB; Bank size: 64MB.
- Open-page row buffer policy.
- First come - first served scheduling policy.
- All the DRAM banks are initially precharged.

(a) [25 points] Identify the memory access pattern (Random/Irregular, Strided (Stride > 1), Stream (Stride = 1)) of each data structure used in the application. Reason about your answer.

source:

dst:


weight:

rank:

old_rank:


(b) [30 points] Consider a case where data structures **rank** and **dst** are mapped in Bank 0, **weight** is mapped in Bank 1, **old_rank** is mapped in Bank 2 and **source** is mapped in Bank 3. Calculate the number of row buffer conflicts after executing App 1 (Lines 9-14) using an input graph with $V=2^{19}$ and $E=2^{22}$. Show your work.

(c) [30 points] Calculate the number of cycles that the processor stalls while waiting for data from memory when executing App 1 using a new input graph with $V=2^{19}$ and $E=2^{22}$ that leads to exactly 2^{20} row buffer conflicts. Show your work.



- (d) [20 points] We replace the DRAM of the previously used system with 2 different DRAM devices, N-DRAM and RL-DRAM.
- Each one of these devices is connected to the processor using a separate channel.
 - N-DRAM has the same characteristics as the DRAM described in the previous questions.
 - RL-DRAM operates with 0.5x the internal latency (i.e. *Precharge, Activate, Read/Write* latencies are halved) of N-DRAM and consists of **only** 1 bank.

Given this system, how would you place the data structures of App 1 across N-DRAM and RL-DRAM to maximize performance? Give your reasoning.



5 Emerging Memory Technologies [105 points]

5.1 Non-Volatile Memory (NVM) [15 points]

Indicate whether each of the following statements is true or false. *Note: we will subtract 1.5 points for each **incorrect** answer. (The minimum score you can get for this question is 0 points.)*

- (a) [3 points] Data is written into Phase Change Memory (PCM) by injecting current to change the magnetic polarity of phase change material.
1. True 2. False
- (b) [3 points] PCM can be denser than DRAM while DRAM is faster and more durable than PCM. Hence, there is an opportunity to benefit from both DRAM and PCM by building a hybrid DRAM-PCM memory system.
1. True 2. False
- (c) [3 points] Multi-Level Cell (MLC) NVM has lower latency and energy consumption than Single-Level Cell (SLC) NVM.
1. True 2. False
- (d) [3 points] NVM has lower endurance than DRAM because writes to NVM take much longer.
1. True 2. False
- (e) [3 points] It takes the same energy to write "00" as to write "11" into an MLC PCM cell.
1. True 2. False

5.2 PCM-based Main Memory [60 points]

A student at ETH wants to build a computer system using PCM as the main memory. Since PCM has limited endurance (i.e., a memory cell fails after 10^7 writes are performed to the cell), the student designs a perfect wear-leveling mechanism (i.e., a mechanism that equally distributes writes across all of memory cells).

The student wants to estimate the worst-case lifetime of PCM when used as main memory. Therefore, she executes a test program to wear out the entire PCM *as quickly as possible*. The test program runs special instructions to bypass the cache hierarchy and repeatedly writes data into different pages until *all* the PCM cells are worn out. The student's measurements show that PCM stops functioning (i.e., all its cells are worn-out) in 3 years. Assume the following:

- The processor employs in-order instruction execution.
- There is no memory-level parallelism (i.e., there is a single bank in the memory system).
- It takes 8 ns to send a memory request from the processor to the memory controller.
- It takes 13 ns to send the request from the memory controller to PCM.
- The write latency of PCM is 43 ns.
- Each write request is fully serialized, i.e., there are three steps of write requests: (1) memory request from CPU to controller, (2) write request from controller to PCM, and (3) data write to PCM cells. None of the steps can be pipelined.
- PCM requests (read or write) are performed at 64-byte granularity.
- PCM adopts a quad-level cell (QLC) technique that stores four bits in a single memory cell.

- (a) [30 points] What is the capacity of the PCM the student uses? Show your work. *Hint:* 3 years $\approx 10^{17}$ ns.

- (b) [30 points] The student decides to improve the lifetime of PCM cells by using the single-level cell (SLC) mode in which a single memory cell stores a single bit. When PCM operates in the SLC mode, each cell's endurance increases by a factor of 10 while the write latency of PCM decreases (compared to the QLC mode). To measure the lifetime of PCM in the SLC mode, the student repeats the same experiment performed in part (a) while keeping everything else in the system the same. The result shows that the lifetime of PCM increases by $2\times$ in SLC mode compared to the QLC mode. Assume the capacity of PCM in QLC mode is 8 GB. What is the write latency of SLC PCM? Show your work.

5.3 Processing using Non-Volatile Memory [30 points]

In class, we learned that a resistive-RAM (ReRAM) crossbar array can significantly accelerate vector-matrix multiplications (VMM). Figure 1 shows the architecture of a 4×4 ReRAM crossbar array designed to perform the following vector-matrix multiplication:

$$I(i_1, i_2, i_3, i_4) \times W \begin{pmatrix} w_{11} & w_{12} & w_{13} & w_{14} \\ w_{21} & w_{22} & w_{23} & w_{24} \\ w_{31} & w_{32} & w_{33} & w_{34} \\ w_{41} & w_{42} & w_{43} & w_{44} \end{pmatrix} = O(o_1, o_2, o_3, o_4).$$

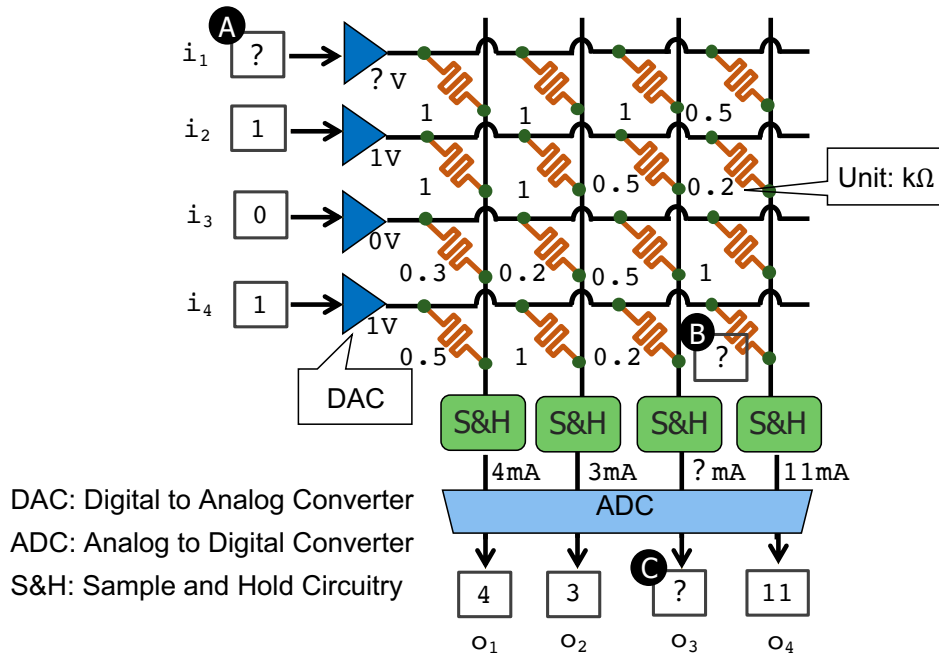


Figure 1: A 4×4 Computation-Capable ReRAM Crossbar Array.

Recall that a typical crossbar array performs a VMM operation based on Kirchhoff’s Law. The VMM computation is conducted via the following four steps.

- Step 1:** We store the matrix into the 2D array of ReRAM cells. Resistance of each ReRAM cell represents the corresponding value in the matrix.
- Step 2:** A DAC (Digital-to-Analog Converter) converts a digital input number to a voltage and applies the voltage to the corresponding wordline.
- Step 3:** A S&H (Sample and Hold) unit collects the current that comes through the corresponding bitline.
- Step 4:** An ADC (Analog-to-Digital Converter) converts the current to the digital number (i.e., an output).

Please identify the unknown values in the three boxes in Figure 1. For your convenience, each number in Figure 1 is converted to a decimal number. Show your work.

- A: _____
- B: _____
- C: _____

6 Asymmetric Multicore [80 points]

A microprocessor manufacturer asks you to design a multicore processor for modern workloads. You should optimize it assuming a workload with 60% of its work in the parallel portion and 40% in the serial portion. You are tasked to compare two configurations that can fit into the processor's die area: 1) Small Cores (SC), and 2) Large + Small Cores (LSC). These consist of the following:

- *SC*: A design that contains 8 small cores, which share the same die. Seven of these small cores operate at a *baseline* fixed instruction throughput. The eighth small core is an overclockable core, which can operate at either: 1) *baseline* throughput, or 2) *overclocked* with $2\times$ the baseline throughput.
- *LSC*: A design that contains 1 large core and 4 small cores that all share the same die. The four small cores operate at the baseline throughput. The large core is $4\times$ faster than a small core.

In addition, Table 1 provides the static power (i.e., when the core is idle) and the dynamic power (i.e., when the core is active) of each of the cores.

Table 1: Power consumption of cores in different modes.

Core	Mode	Static Power (W)	Dynamic Power (W)
Small	Baseline	0.5	1
	Overclocked	1	8
Large	Baseline	2	4

The SC processor executes the parallel portion on *all* the small cores (including the overclockable core, operating at the baseline throughput), and the serial portion *only* on the overclockable core (using either baseline or overclocked options). The LSC processor executes the parallel portion only on the small cores, and the serial portion only on the large core.

Please answer the following questions.

- (a) [20 points] Which of the three design configurations (SC, SC with overclocked core, or LSC) results in the highest performance? Show your work.

- (b) [20 points] The energy consumption should also be a metric of reference in your design. Which of the three design configurations (SC, SC with overclocked core, or LSC) results in the lowest energy consumption? Show your work.

- (c) [20 points] At least what ratio of a workload should be spent on the parallel section so that the *SC* configuration, *even without overclocking*, performs better than the *LSC* configuration? Show your work.

- (d) [20 points] In order to improve the performance of the LSC configuration, you come up with hardware design optimizations to improve the throughput of the large core. You expect that these optimizations will increase the throughput of the large core by $T\times$. Given an application with 90% of its work in the parallel portion, is it possible for the LSC configuration to outperform the SC configuration *with* overclocking? If yes, for which values of T ? Show your work.

7 Interconnection Networks [115 points]

Suppose you want to connect 16 processors using two topologies: 4×4 Mesh and 4×4 Torus with bi-directional links, similar to Figure 1:

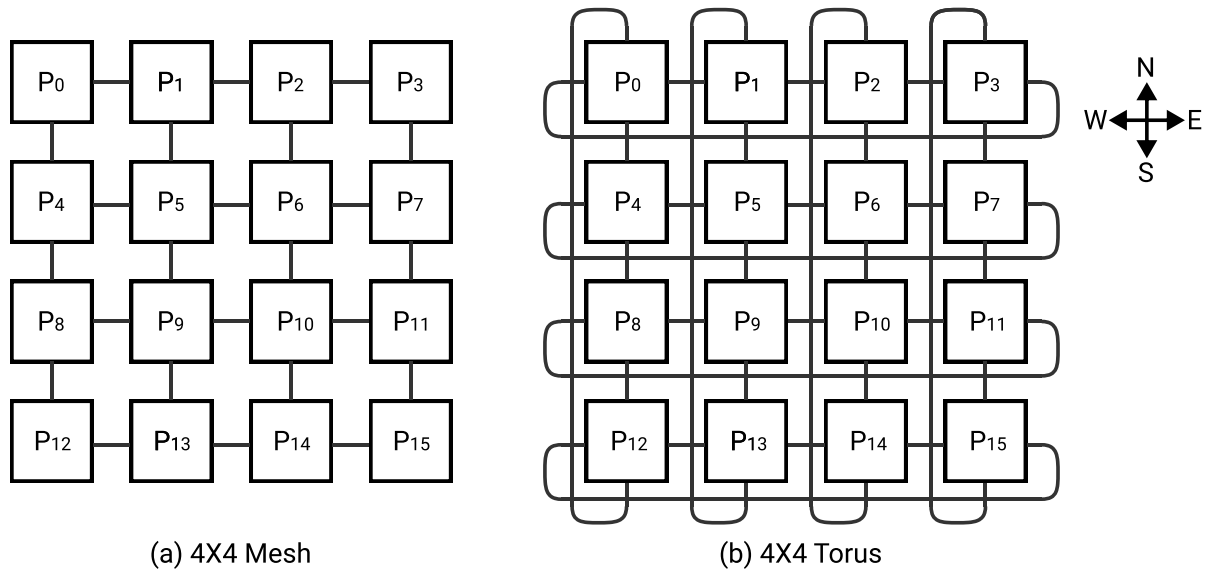


Figure 1: 2D mesh and 2D torus topologies connecting 16 processors.

You know that there are eight possible turns in the 2D mesh and the 2D torus topologies, as shown in Figure 2:




Figure 2: Eight possible turns in 2D mesh and 2D torus topologies.

You are considering three different *minimal* routing algorithms based on which turns are forbidden:

- **Routing Algorithm 1:** Turns NE, SW, NW, and SE are forbidden. The other four turns are allowed.
- **Routing Algorithm 2:** Turns SW and NW are forbidden. The other six turns are allowed.
- **Routing Algorithm 3:** Turns WN and NW are forbidden. The other six turns are allowed.

For the rest of this question, we address source-destination pairs using $\langle \text{src}, \text{dest} \rangle$ representation. Please answer the following questions and show your work.

- (a) [50 points] Show all possible paths between source and destination pairs $\langle P_5, P_{11} \rangle$ and $\langle P_5, P_{12} \rangle$, using each of the previously mentioned routing algorithms. For each path, mention the source router, intermediate routers, and destination router.

2D Mesh

Algorithm 1, $\langle P_5, P_{11} \rangle$

Algorithm 1, $\langle P_5, P_{12} \rangle$

Algorithm 2, $\langle P_5, P_{11} \rangle$

Algorithm 2, $\langle P_5, P_{12} \rangle$

Algorithm 3, $\langle P_5, P_{11} \rangle$

Algorithm 3, $\langle P_5, P_{12} \rangle$

2D Torus

Algorithm 1, $\langle P_5, P_{11} \rangle$

Algorithm 1, $\langle P_5, P_{12} \rangle$




Algorithm 2, $\langle P_5, P_{11} \rangle$

Algorithm 2, $\langle P_5, P_{12} \rangle$

Algorithm 3, $\langle P_5, P_{11} \rangle$


Algorithm 3, $\langle P_5, P_{12} \rangle$

- (b) [45 points] Which of these three algorithms are deadlock free? For each case, if deadlocks can happen, provide a deadlock scenario. Otherwise, prove that deadlock occurrence is impossible (Hint: Use contradiction).


2D Mesh**Routing Algorithm 1:****Routing Algorithm 2:****Routing Algorithm 3:**

2D Torus

Routing Algorithm 1:



Routing Algorithm 2:




Routing Algorithm 3:




You are about to design a routing algorithm for each of these two topologies with the following two requirements: Your routing algorithm (1) should be deadlock free, and (2) should *not* be deterministic.

- (c) [20 points] Which routing algorithm out of the three earlier routing algorithms do you choose for each of these two topologies? Explain your reasoning.

2D mesh



2D torus



8 BONUS: Prefetching using Reinforcement Learning [70 points]

You are designing a hardware prefetcher for a processor using a reinforcement learning (RL) agent, as discussed in lecture about the Pythia prefetcher. For a memory request to cacheline address A , the prefetcher selects a prefetch offset O and issues a prefetch memory request to cacheline address $A + O$. For every prefetch request, the memory hierarchy provides a numerical reward R to the prefetcher that can take a value of one of the following three reward levels:

- Accurate (R_A), signifying that the prefetch request was demanded by the processor.
- Inaccurate (R_{IN}), signifying that the prefetch request was not demanded by the processor.
- No-prefetch (R_{NP}), signifying that the prefetcher did not prefetch anything.

In the initial configuration of the prefetcher, you set the values of the reward levels as follows: $R_A = 20$, $R_{IN} = -4$, and $R_{NP} = -10$.

Recall that the **coverage** of a prefetcher is defined as the fraction of a program's memory requests correctly prefetched by the prefetcher, while the **accuracy** of a prefetcher is defined as the fraction of prefetched requests that are actually demanded by the program.

- (a) [25 points] Which of the following statements, if any, are **CORRECT** if you set $R_{NP} = 1000$ in the initial prefetcher configuration? All other reward level values, except R_{NP} , remain the same as the initial configuration. Select **ALL** that apply and explain briefly. You may get partial credits for a partially-complete answer, given a correct explanation.

- A. The coverage of the prefetcher will significantly *increase*.
- B. The coverage of the prefetcher will significantly *decrease*.
- C. The prefetcher will start prefetching aggressively.
- D. The accuracy of the prefetcher may *increase*.

- (b) [25 points] Which of the following statements, if any, are **CORRECT** if you set $R_{IN} = -500$ in the initial prefetcher configuration? All other reward level values, except R_{IN} , remain the same as the initial configuration. Select **ALL** that apply and explain briefly. You may get partial credits for a partially-complete answer, given a correct explanation.

- A. The prefetcher will aggressively prefetch, even though the prefetches might be inaccurate.
- B. The accuracy of the prefetcher will likely *increase*.
- C. The accuracy of the prefetcher will likely *decrease*.
- D. The coverage of the prefetcher might *increase or decrease*.

- (c) [20 points] You want to make the prefetcher system-aware by incorporating into the decision making process *the type of power source* used in the system. You want to configure the prefetcher in the following way:
- i. The prefetcher should generate accurate prefetches whenever possible, irrespective of whether the system is connected to an external power adapter or running on a battery.
 - ii. If the system is connected to an external power adapter, the prefetcher should continue to prefetch, even if the prefetch might be inaccurate.
 - iii. If the system is running on battery, the prefetcher should prefer not to generate a prefetch request if the prefetch is likely to be inaccurate.

How would you configure the values of the three reward levels? The relative ordering of the reward level values is sufficient for a correct answer, rather than their exact values. Fill in the blanks below and explain your reasoning briefly in the provided box.

```
if (power_source == battery)
```

```
    [ ] < [ ] < [ ]
```

```
else if (power_source == external_adapter)
```

```
    [ ] < [ ] < [ ]
```