

ETH 263-2210-00L COMPUTER ARCHITECTURE, FALL 2021

HW 1: DRAM REFRESH

Instructor: Prof. Onur Mutlu

TAs: Juan Gómez Luna, Mohammed Alser, Jisung Park, Lois Orosa Nogueira, Gagandeep Singh, Haiyu Mao, Behzad Salami, Nour Almadhoun Alserr, Mohammad Sadr, Hasan Hassan, Can Firtina, Geraldo Francisco De Oliveira Junior, Abdullah Giray Yaglikci, Rahul Bera, Konstantinos Kanellopoulos, Nika Mansouri Ghiasi, Rakesh Nadig, João Dinis Ferreira, Haocong Luo, Roknoddin Azizibarzoki

Given: Friday, Oct 05, 2021

Due: **Friday, Oct 19, 2021**

- **Handin - Critical Paper Reviews (1).** You need to submit your reviews to <https://safari.ethz.ch/review/architecture21/>. Please, check your inbox, you should have received an email with the password you should use to login. If you did not receive any email, contact comparch@lists.inf.ethz.ch. In the first page after login, you should click in "Computer Architecture Home", and then go to "any submitted paper" to see the list of papers.
- **Handin - Questions (2-7).** You should upload your answers to the Moodle Platform (<https://moodle-app2.1et.ethz.ch/course/view.php?id=15536>) as a single PDF file.

1. Critical Paper Reviews [1,000 points]

Please read the guidelines for reviewing papers and check the sample reviews. We also assign you three **required readings** for this homework. You may access them by *simply clicking on the QR codes below or scanning them*. We will give out extra credit that is worth 0.5% of your total grade for each good review. If you review a paper other than the REQUIRED papers, you will receive 250 BONUS points on top of 1,000 points you may get from paper reviews (i.e., each additional submission is worth 250 BONUS points with a possibility to get up to 3,000 points).



Guidelines



Sample reviews



Required Reading 1



Required Reading 2



Required Reading 3

Write an approximately one-page critical review for the following required readings (i.e., papers from #1 to #3) **and at least 1** of the remaining 9 papers (i.e., papers from #4 to #12). A review with bullet point style is more appreciated. Try not to use very long sentences and paragraphs. Keep your writing and sentences simple. Make your points bullet by bullet, as much as possible.

1. **(REQUIRED)** Richard Hamming, "You and Your Research," Bell Communications Research Colloquium Seminar, 1986. <https://safari.ethz.ch/architecture/fall2021/lib/exe/fetch.php?media=youandyourresearch.pdf>
2. **(REQUIRED)** Onur Mutlu, "Intelligent Architectures for Intelligent Machines," World Conference on Energy Science and Technology (WCEST), 2021.
 - Talk: https://www.youtube.com/watch?v=20Bdcw-ilQY&list=PL5Q2soXY2Zi8D_5MGV6EnXEJHnV2YFBJl&index=61
 - Slide (PPT): <https://people.inf.ethz.ch/omutlu/pub/onur-TUBA-WCEST-Keynote-IntelligentArchitecturesForIntelligentSystems-August-11-2021.pptx>
 - Slide (PDF): <https://people.inf.ethz.ch/omutlu/pub/onur-TUBA-WCEST-Keynote-IntelligentArchitecturesForIntelligentSystems-August-11-2021.pdf>
3. **(REQUIRED)** Onur Mutlu, "Main Memory Scaling: Challenges and Solution Directions," Invited Book Chapter in More than Moore Technologies for Next Generation Computer Design, pp. 127-153, Springer, 2015. https://people.inf.ethz.ch/omutlu/pub/main-memory-scaling_springer15.pdf
4. Thomas Moscibroda and Onur Mutlu, "Memory Performance Attacks: Denial of Memory Service in Multi-Core Systems," in Proceedings of the USENIX Security Symposium, 2007. https://people.inf.ethz.ch/omutlu/pub/mph_usenix_security07.pdf
5. Liu et al., "RAIDR: Retention-Aware Intelligent DRAM Refresh," in Proceedings of the International Symposium on Computer Architecture, 2012. https://people.inf.ethz.ch/omutlu/pub/raidr-dram-refresh_isca12.pdf
6. Liu et al., "An Experimental Study of Data Retention Behavior in Modern DRAM Devices: Implications for Retention Time Profiling Mechanisms," in Proceedings of the International Symposium on Computer Architecture, 2013. https://people.inf.ethz.ch/omutlu/pub/dram-retention-time-characterization_isca13.pdf
7. Kim et al., "Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors," in Proceedings of the International Symposium on Computer Architecture, 2014. https://people.inf.ethz.ch/omutlu/pub/dram-row-hammer_isca14.pdf
8. Kim et al., "Revisiting RowHammer: An Experimental Analysis of Modern Devices and Mitigation Techniques," in Proceedings of the International Symposium on Computer Architecture, 2020. https://people.inf.ethz.ch/omutlu/pub/Revisiting-RowHammer_isca20.pdf
9. Frigo et al., "TRRespass: Exploiting the Many Sides of Target Row Refresh," in Proceedings of the IEEE Symposium on Security and Privacy, 2020. https://people.inf.ethz.ch/omutlu/pub/rowhammer-TRRespass_ieee_security_privacy20.pdf
10. Burton H. Bloom, "Space/Time Trade-Offs in Hash Coding with Allowable Errors," Communications of the ACM, 1970. https://safari.ethz.ch/architecture/fall2021/lib/exe/fetch.php?media=bloom_commacm_1970.pdf
11. Venkatesan et al., "Retention-Aware Placement in DRAM (RAPID): Software Methods for Quasi-Non-Volatile DRAM," in Proceedings of the International Symposium on High-Performance Computer Architecture, 2006. https://safari.ethz.ch/architecture/fall2021/lib/exe/fetch.php?media=venkatesan_hpca_2006.pdf
12. Chang et al., "Improving DRAM Performance by Parallelizing Refreshes with Accesses," in Proceedings of the International Symposium on High-Performance Computer Architecture, 2014. https://people.inf.ethz.ch/omutlu/pub/dram-access-refresh-parallelization_hpca14.pdf

2. Main Memory [150 points]

Answer the following questions for a machine that has a 4 GB DRAM main memory system and each row is refreshed every 64 ms.

- (a) During standalone runs of two applications (A and B) on the machine, you observe that application A spends a significantly larger fraction of cycles stalling for memory than application B does while both applications have a similar number of memory requests. What might be the reasons for this?

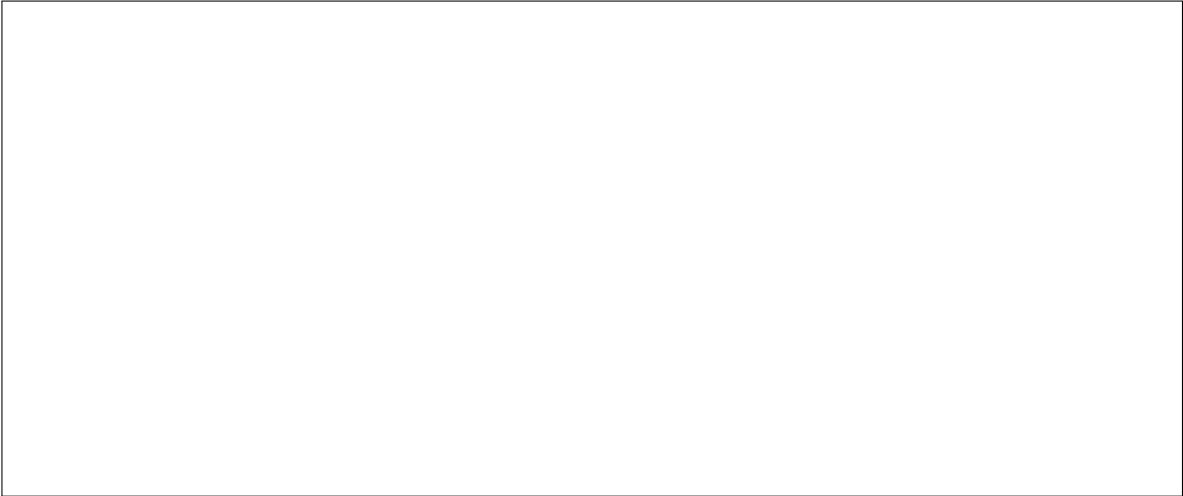
- (b) Application A also consumes a much larger amount of memory energy than application B does. What might be the reasons for this?

- (c) When applications A and B run together on the machine, application A's performance significantly degrades, while application B's performance does not degrade as much. Why might this happen?

- (d) The designer decides to use a smarter policy to refresh the memory. A row is refreshed only if it has not been accessed in the past 64 ms. Do you think this is a good idea? Why or why not?



- (e) When this new refresh policy is applied, the refresh energy consumption drops significantly during a run of application B. In contrast, during a run of application A, the refresh energy consumption reduces only slightly. Why might this happen?



3. DRAM Refresh - Utilization [150 points]

A memory system has four channels, and each channel has two ranks of DRAM chips. Each memory channel is controlled by a separate memory controller. Each rank of DRAM contains eight banks. A bank contains 32K rows. Each row in one bank is 8KB. The minimum retention time among all DRAM rows in the system is 64 ms. In order to ensure that no data is lost, every DRAM row is refreshed once per 64 ms. Every DRAM row refresh is initiated by a command from the memory controller which occupies the command bus on the associated memory channel for 5 ns and the associated bank for 40 ns. Let us consider a 1.024 second span of time.

We define *utilization* (of a resource such as a bus or a memory bank) as the fraction of total time for which a resource is occupied by a refresh command.

For each calculation in this section, you may leave your answer in *simplified* form in terms of powers of 2 and powers of 10.

- (a) How many refreshes are performed by the memory controllers during the 1.024 second period in total across all four memory channels?

- (b) What command bus utilization, across all memory channels, is directly caused by DRAM refreshes?

- (c) What data bus utilization, across all memory channels, is directly caused by DRAM refreshes?

- (d) What bank utilization (on average across all banks) is directly caused by DRAM refreshes?

- (e) The system designer wishes to reduce the overhead of DRAM refreshes in order to improve system performance and reduce the energy spent in DRAM. A key observation is that not all rows in the DRAM chips need to be refreshed every 64 ms. In fact, rows need to be refreshed only at the following intervals in this particular system:

Required Refresh Rate	Number of Rows
64 ms	2^5
128 ms	2^9
256 ms	all other rows

Given this distribution, if all rows are refreshed only as frequently as required to maintain their data, how many refreshes are performed by the memory controllers during the 1.024 second period in total across all four memory channels?

What command bus utilization (as a fraction of total time) is caused by DRAM refreshes in this case?

(f) What DRAM data bus utilization is caused by DRAM refreshes in this case?

(g) What bank utilization (on average across all banks) is caused by DRAM refreshes in this case?

(h) The system designer wants to achieve this reduction in refresh overhead by refreshing rows less frequently when they need less frequent refreshes. In order to implement this improvement, the system needs to track every row's required refresh rate. What is the minimum number of bits of storage required to track this information?

(i) Assume that the system designer implements an approximate mechanism to reduce refresh rate using Bloom filters, as we discussed in class. One Bloom filter is used to represent the set of all rows which require a 64 ms refresh rate, and another Bloom filter is used to track rows which require a 128 ms refresh rate. The system designer modifies the memory controller's refresh logic so that on every potential refresh of a row (every 64 ms), it probes both Bloom filters. If either of the Bloom filter probes results in a "hit" for the row address, and if the row has not been refreshed in the most recent length of time for the refresh rate associated with that Bloom filter, then the row is refreshed. (If a row address hits in both Bloom filters, the more frequent refresh rate wins.) Any row that does not hit in either Bloom filter is refreshed at the default rate of once per 256 ms.

The false-positive rates for the two Bloom filters are as follows:

Refresh Rate Bin	False Positive Rate
64 ms	2^{-20}
128 ms	2^{-8}

The distribution of required row refresh rates specified in part (e) still applies.

How many refreshes are performed by the memory controllers during the 1.024 second period in total across all four memory channels?

What command bus utilization results from this refresh scheme?

What data bus utilization results from this refresh scheme?

What bank utilization (on average across all banks) results from this refresh scheme?

4. VRL: Variable Refresh Latency (VRL) [150 points]

In this question, you are asked to evaluate "Variable Refresh Latency," proposed by Das et al. in the Design Automation Conference 2018.¹

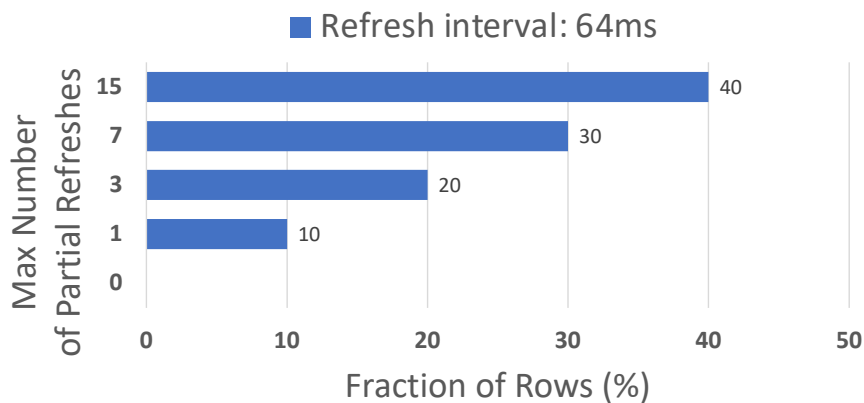
The paper presents two key observations:

- First, a cell's charge reaches 95% of the maximum charge level in 60% of the nominal latency value during a refresh operation. In other words, the last 40% of the refresh latency is spent to increase the charge of a cell from 95% to 100%. Based on this observation, the paper defines two types of refresh operations: (1) *full refresh* and (2) *partial refresh*. Full refresh uses the nominal latency value and restores the cell charge to 100%, while the latency of partial refresh is only 60% of the nominal latency value and it restores 95% of the charge.
- Second, a fully refreshed cell operates correctly even after multiple partial refreshes, but it needs to be fully refreshed again after a finite number of partial refreshes. The maximum number of partial refreshes before a full refresh is required varies from cell to cell.

The **key idea** of the paper is to apply a *full refresh* operation **only when necessary** and use *partial refresh* operations at all other times.

(a) Consider a case in which:

- Each row must be refreshed every 64 ms. In other words, the refresh interval is 64 ms.
- Row refresh commands are evenly distributed across the refresh interval. In other words, all rows are refreshed exactly once in any given 64 ms time window.
- You are given the following plot, which shows *the distribution of the maximum number of partial refreshes* across all rows of a particular bank. For example, if the maximum number of refreshes is three, those rows can be partially refreshed for at most three refresh intervals, and the fourth refresh operation must be a full refresh.
- If all rows were always fully refreshed, the time that a bank is busy serving the refresh requests within a refresh interval would be T .



How much time does it take (in terms of T) for a bank to refresh all rows within a refresh interval, after applying Variable Refresh Latency?

¹Das, A. et al., "VRL-DRAM: Improving DRAM Performance via Variable Refresh Latency." In Proceedings of the 55th Annual Design Automation Conference (DAC), 2018.



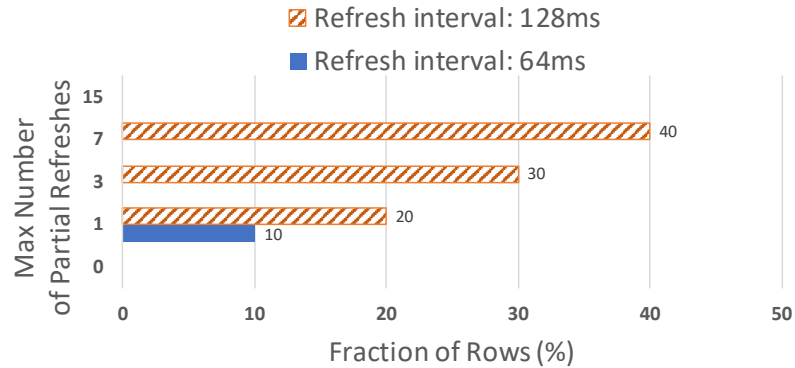
- (b) You find out that you can relax the refresh interval, and define your baseline as follows:
- 75% of the rows are refreshed at every 128ms; 25% of the rows are refreshed at every 64ms.
 - Refresh commands are evenly distributed in time.
 - All rows are always fully refreshed.
 - A single refresh command costs $0.2/N$ ms, where N is the number of rows in a bank.
 - *Refresh overhead* is defined as the fraction of time that a bank is busy, serving the refresh requests over a very large period of time.

Calculate the refresh overhead for the baseline.



(c) Consider a case where:

- 90% of the rows are refreshed at every 128ms; 10% of the rows are refreshed at every 64ms.
- Refresh commands are evenly distributed in time.
- You are given the following plot, which shows *the distribution of the maximum number of partial refreshes* across all rows of a particular bank.
- A single refresh command costs $0.2/N$ ms, where N is the number of rows in a bank.
- *Refresh overhead* is defined as the fraction of time that a bank is busy, serving the refresh requests over a very large period of time.



Calculate the refresh overhead. Show your work step-by-step. Then, compare it against the baseline configuration (part b). How much reduction ($1 - \frac{new}{old}$) do you see in the performance overhead of refreshes?

5. RowHammer [150 points]

5.1. RowHammer Attacks

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

- The CPU has a single 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 open-page policy.
- The DRAM modules and the memory controller do not employ any remapping or scrambling schemes for the physical address.
- All the cells in the DRAM subsystem are equally vulnerable to RowHammer-induced errors.

You implement codes based on 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) You run Code 1 below, but you cannot observe any errors in the target system. You figured out that the number of activations is much lower than your expectation. Give reason(s) as to why Code 1 cannot introduce a sufficient amount of activations.

Code 1

```
1: LOOP:
2:   STORE 0x8732, R0
3:   CLFLUSH 0x8732
4:   B LOOP
```

(b) You try Codes 2a, 2b, and 2c, but find that *only one of them can induce RowHammer errors* in your DRAM subsystem. Which code segment is the one that can induce RowHammer errors? Justify your answer.

Code 2a

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

Code 2b

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

Code 2c

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

5.2. RowHammer Mitigation Mechanisms

To identify a viable RowHammer mitigation mechanism for your system, you compare the two following mitigation mechanisms:

- **Mechanism A.** The memory controller maintains a counter for every row, which increments every time the corresponding row is activated. If the counter value for a row exceeds a threshold value T , the memory controller activates the row's two adjacent rows and resets the counter.
- **Mechanism B.** Each time a row is closed (or precharged), the memory controller flips a biased coin with a probability p of turning up heads, where $p \ll 1$. If the coin turns up heads, the memory controller activates one of its adjacent rows where either of the two adjacent rows are selected with equal probability ($p/2$).

- (a) You set T for Mechanism A to 164 K based on the value of the Maximum Activation Count (MAC, i.e., the maximum number of times a row can be activated without inducing RowHammer errors in its adjacent rows) reported by the DRAM manufacturer. Calculate the number of bits required for counters in a memory controller which manages a single channel, 2 ranks per channel, 8 banks per rank, and 2^{15} rows per bank.

- (b) How does the answer to (a) change when both the number of rows per bank and the number of banks per chip are doubled?

- (c) You profile the memory access pattern of the target system, and observe that the same pattern repeats exactly every 64 ms (the current refresh interval). Table 2 shows the number of activations for each row within a 64-ms time interval in a descending order. Given values $T = 164$ K for Mechanism A and $p = 0.001$ for Mechanism B, calculate the expected number of additional activations within a 64-ms time interval under each technique.

Row Index	# of ACTs
0x7332F	73 K
0x1802C	64 K
0x03F05	32 K
0x5FF02	10 K
...	...
Total	480 K

Table 2. Number of Activations for Each Row.

(d) How does the answer to (c) change when both the number of rows per bank and the number of banks per chip are doubled? Assume that the memory access pattern does not change.

(e) What is the *common challenge* to implement the above mechanisms in the commodity systems?

(f) How can you address the common challenge?

6. RowHammer [150 points]

6.1. RowHammer Properties

Determine 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 point.)*

- (a) [3 points] Violating DRAM timing parameters is necessary to induce RowHammer bit flips.

1. True 2. False

- (b) [3 points] SECDED (Single Error Correction Double Error Detection) Hamming ECC cannot guarantee RowHammer-safe operation.

1. True 2. False

- (c) [3 points] Mobile devices are RowHammer-safe because they use low power memory chips.

1. True 2. False

- (d) [3 points] We can more effectively induce bit flips in a given victim row by hammering rows in different banks.

1. True 2. False

- (e) [3 points] In DRAM TRR (Target Row Refresh) mechanism provides RowHammer-safe operation. The only problem is that it is not implemented in all DRAM chips.

1. True 2. False

6.2. RowHammer Mitigation

You are assigned to implement a RowHammer mitigation mechanism within the memory controller of a new processor. The DRAM chips that will be used with the processor are organized as 1 channel, 1 rank, 8 banks, 8 KB row size, and 1 GB total capacity. The DRAM protocol specifies that each row should be refreshed every 64 ms, and there should be at least 64 ns between two row activations targeting the same DRAM bank. The RowHammer threshold of the chips is 50,000 row activations per aggressor row during a double-sided attack, while no hammer count for a single-sided attack can induce bit flips within 64 ms.

- (a) How many rows in a bank can an attacker concurrently hammer enough times to induce bit flips? Show your work and explain clearly.

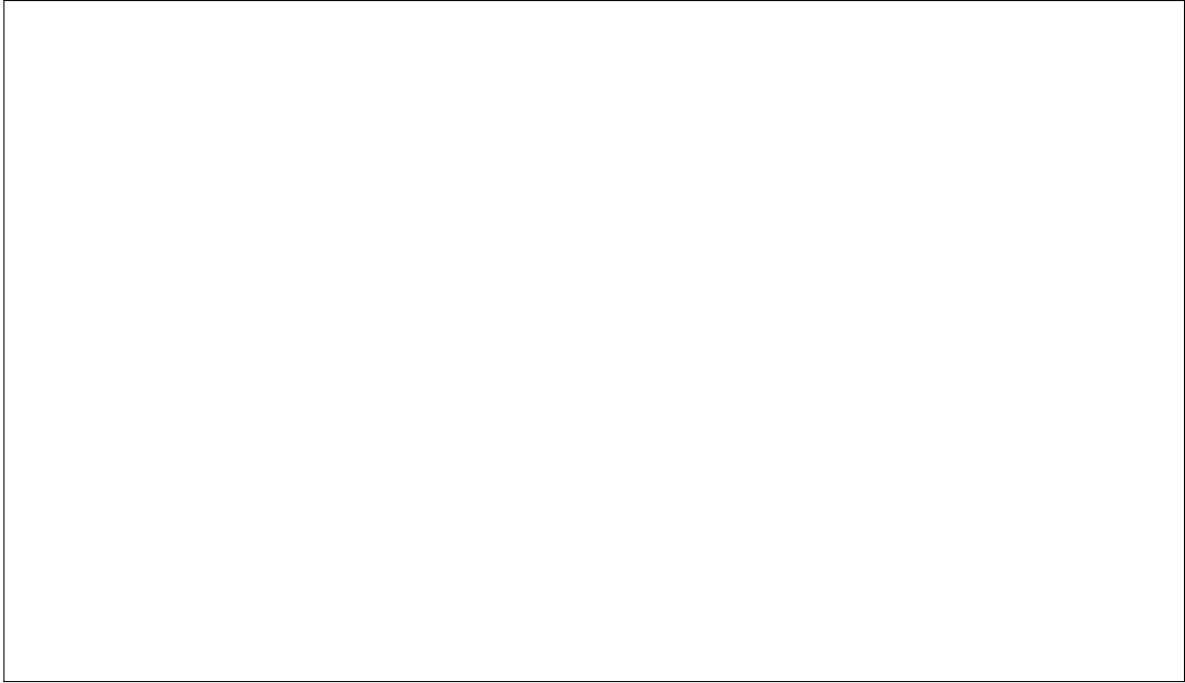
(b) What is the maximum number of victim rows that can be affected in a bank? Show your work and explain clearly.

(c) To detect a RowHammer attack, you consider using existing hardware performance counters. Sort the following performance metrics from *the most accurate to the least accurate* as an indicator of an ongoing RowHammer attack. Explain your reasoning clearly. You may not get points without explanation even if the ordering is correct.

- AM: Available Memory: The amount of physical memory, free and available to processes.
- RBHR: Row Buffer Hit Rate: The fraction of requests that hit the row buffer
- LLC MPKI: Last Level Cache Misses per Kilo Instructions
- RBCPKI: Row Buffer Conflicts per Kilo Instructions

(d) As the RowHammer mitigation mechanism's designer, you have a chance to change the memory controller's row policy. From the perspective of RowHammer mitigation, which page policy would you choose? Please mark below and explain.

Open-Row Policy Closed-Row Policy Does Not Matter



7. BONUS: DRAM Refresh - Energy [150 points]

A new supercomputer has a DRAM-based memory system with the following configuration:

- The total capacity is 1 ExaByte (EB).
- The DRAM row size is 8 KiloByte (KB).
- The minimum retention time among all DRAM rows in the system is 64 ms. In order to ensure that no data is lost, every DRAM row is refreshed once every 64 ms. (Note: For each calculation in this question, you may leave your answer in simplified form in terms of powers of 2 and powers of 10.)

(a) How many DRAM rows does the memory system have?

(b) According to the specs of this hypothetical DRAM device, the memory controller should issue a refresh command (REF) at every 7.8 μ s. How many rows is refreshed when a REF command is issued?

(c) What is the power consumption of DRAM refresh? (Hint: you will need to figure out how much current the DRAM device draws during refresh operations. You can find useful information in the technical note by Micron.² Use the current (IDD) numbers specified in Micron's datasheet.³ Clearly state all the assumptions and show how you derive the power numbers. You are welcome to use other datasheets as well. Make sure you specify how you obtain the power numbers and show your calculations and thought process.)

(d) What is the total energy consumption of DRAM refresh during a refresh cycle? And during a day?

²Micron, "TN-47-04: Calculating Memory System Power for DDR2 Introduction"

³Micron, "1Gb: x4, x8, x16 DDR3 SDRAM Features"