

*ZORO:
Zack's Open-Row
Oriented Memory
Scheduler*

Jack Davidson and
Zackery Painter

Introduction

- The goal was to determine if we can make a better scheduling algorithm using a predictor
- Originally wanted to predict if we should leave the page open or close the page after a read / write
- However, decided it would be more appropriate to predict what we should send from the CPU to the memory controller

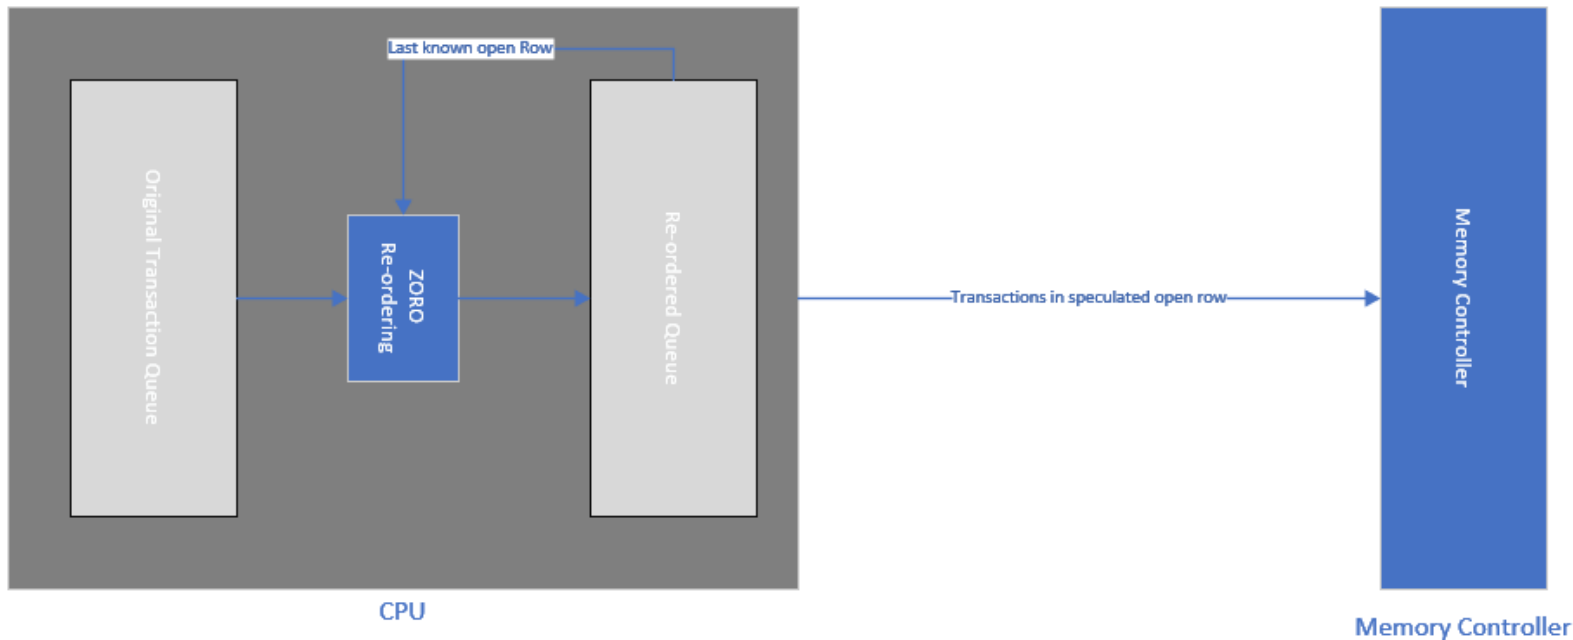


Figure 1 - General idea of ZORO

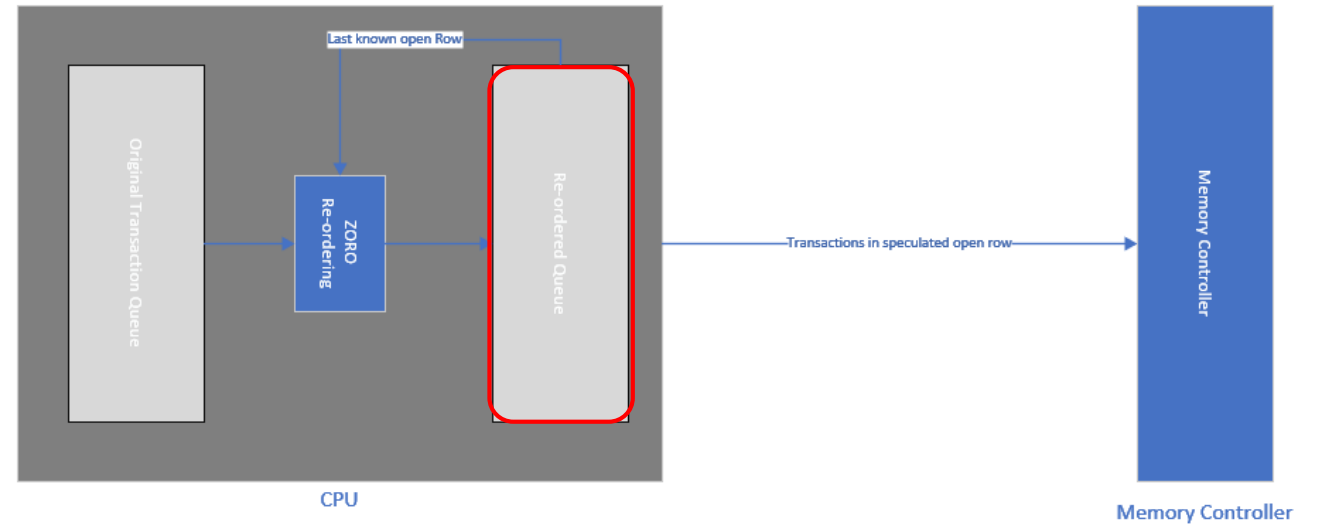
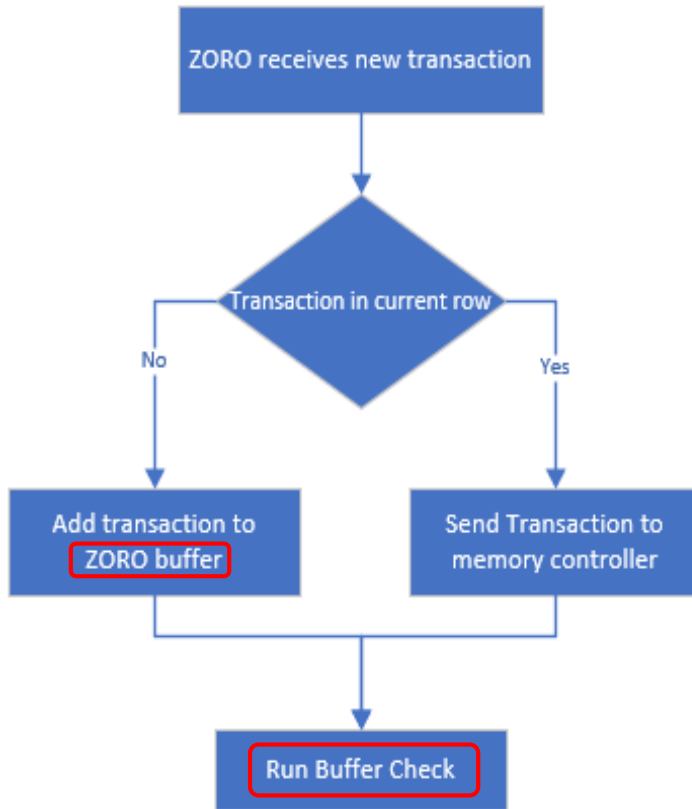
ZORO

- Important vocabulary

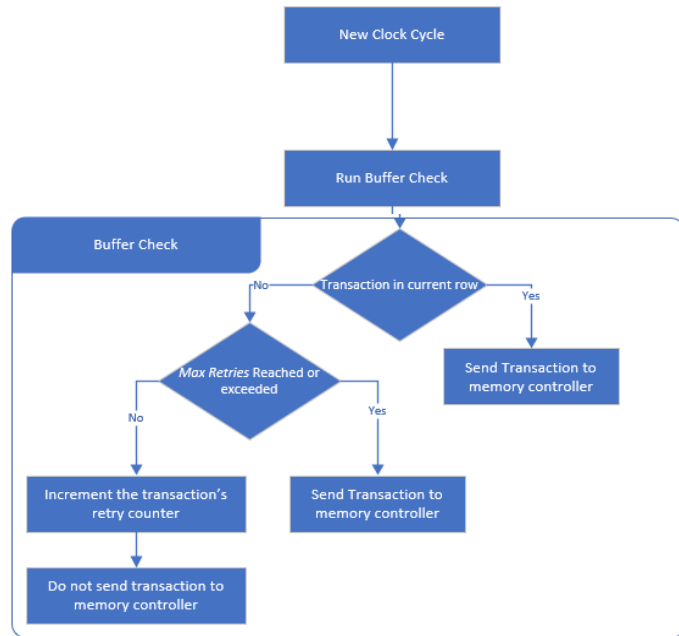
- *Transaction Groups* - A transaction group is group of memory transactions that are in the same row

Transaction	Address	Row
T1	0x0A	1
T2	0xAB	1

- *Transaction that transactions go before sent to the memory controller*



ZORO



- How does ZORO work?
 - *ZORO aims to take advantage of FR-FCFS by always attempting to send groups of memory transactions that are in the current row*
- *When a new transaction is received:*
 - Calculate transaction's row
 - Compare to last known opened row
 - If equal, send to memory controller
 - Otherwise, put in ZORO buffer
- Run a buffer check
- On each clock cycle
 - Run a buffer check

Simulations

- Gem5 with a side of DRAMsim3
- Put it in the oven for 100 iterations
- Gives a good balance of command flavor and runtime spice
- We stuck to the standard Coremark workload, didn't have time to experiment with other benchmark recipes
- Tried a few different cache sizes to increase the number of memory transactions and expand out palates



Simulations

- Why DRAMSim3?
 - Interfaced “nicely” with GEM5
 - Very easy to change and look at source code
 - More accurate than GEM5’s DRAM modules
 - Very short compile time
- Why GEM5?
 - Very robust and easy to create workloads for
 - Very well documented on their website
 - Optimized compiler and well commented code



Simulations

(Additional details)

- DRAM: DDR4_8Gb_x8_2400
- CPU: O3CPU (Basic Out-of-Order CPU)
- Default Cache size:
 - *L1d - 64kB*
 - *L1i - 32kB*
 - *L2 - 2MB*
 - *L3 - 16MB*



We got...

Data

(lots of data)

15 data files x 10 folders = 150 data files!

(Not even including testing data and text files!)

The image is a collage of data-related elements. On the left, a file explorer shows a directory structure with files named 'dramsims' in JSON and TXT formats. A code snippet in the center shows a Python loop that reads energy data from files and processes it. To the right, a cartoon illustration features a bottle labeled 'SPICY DATA' with a flame icon, positioned above a bar chart with five bars of varying heights and colors (pink, cyan, orange, green, red). The background is a dark grey grid.

File Name	File Type	Size
dramsims3_1.json	JSON File	12 KB
dramsims3_1.txt	Text Document	6 KB
dramsims3_2.json	JSON File	12 KB
dramsims3_2.txt	Text Document	6 KB
dramsims3_3.json	JSON File	12 KB
dramsims3_3.txt	Text Document	6 KB
dramsims3_4.json	JSON File	12 KB
dramsims3_4.txt	Text Document	6 KB
dramsims3_11.json	JSON File	13 KB
dramsims3_11.txt	Text Document	6 KB
dramsims3_12.json	JSON File	13 KB
dramsims3_12.txt	Text Document	6 KB
dramsims3_13.json	JSON File	13 KB
dramsims3_13.txt	Text Document	6 KB
dramsims3_14.json	JSON File	12 KB


```
lookFor="read_energy"  
df=pd.DataFrame()  
for i in range(0, len(  
    exec("df"+str(i)+  
    exec("df"+str(i)+  
    r = 16/int(caches  
    if(r<1):  
  
        r = r*1024  
    exec("df['Default
```


Results

Overview:

- Row hit rate did what we expected when cache was normal, less so for the smaller configs
- Energy also did what we expected for the initial setup
- The Baseline hit rate did some strange things, crossing over read/write as cache decreases

Activation Energy Results (Default cache)



- The left shows our activation energy using ZORO as we change the max retries variable
- Notice an interesting trend!

Our best value at Max Retries = 3 was 2.93%!

Blue line - ZORO results
Orange line - Baseline results (unmodified)

Average Read Latency Results (Default Cache)


Average Read Latency vs Max_Retries



- Shows the overall read latency plotted against the ZORO max retries parameter
- Overall, ZORO seems to perform better for every value of max retries than the baseline

Total Energy (Default Cache)

- Overall, less total energy than Baseline!




Memory Row Hit Rate Improvement Results (Default Cache)

Memory Row Hit Rate Improvement vs Max Retries at Default Cache Size

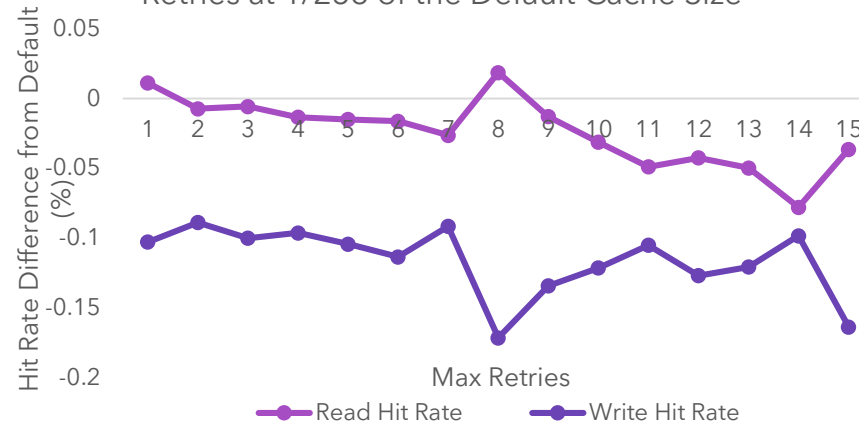


What about smaller cache sizes?

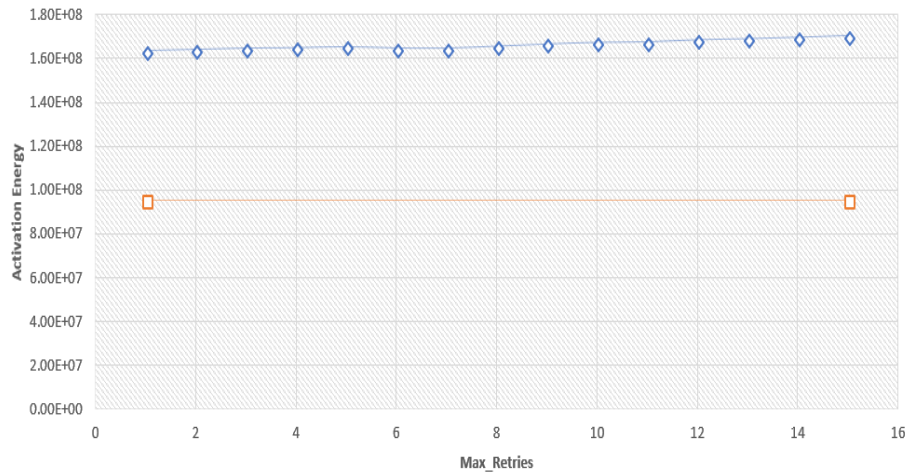
Memory Row Hit Rate Improvement vs Max Retries at 1/2 of the Default Cache Size




Memory Row Hit Rate Improvement vs Max Retries at 1/256 of the Default Cache Size

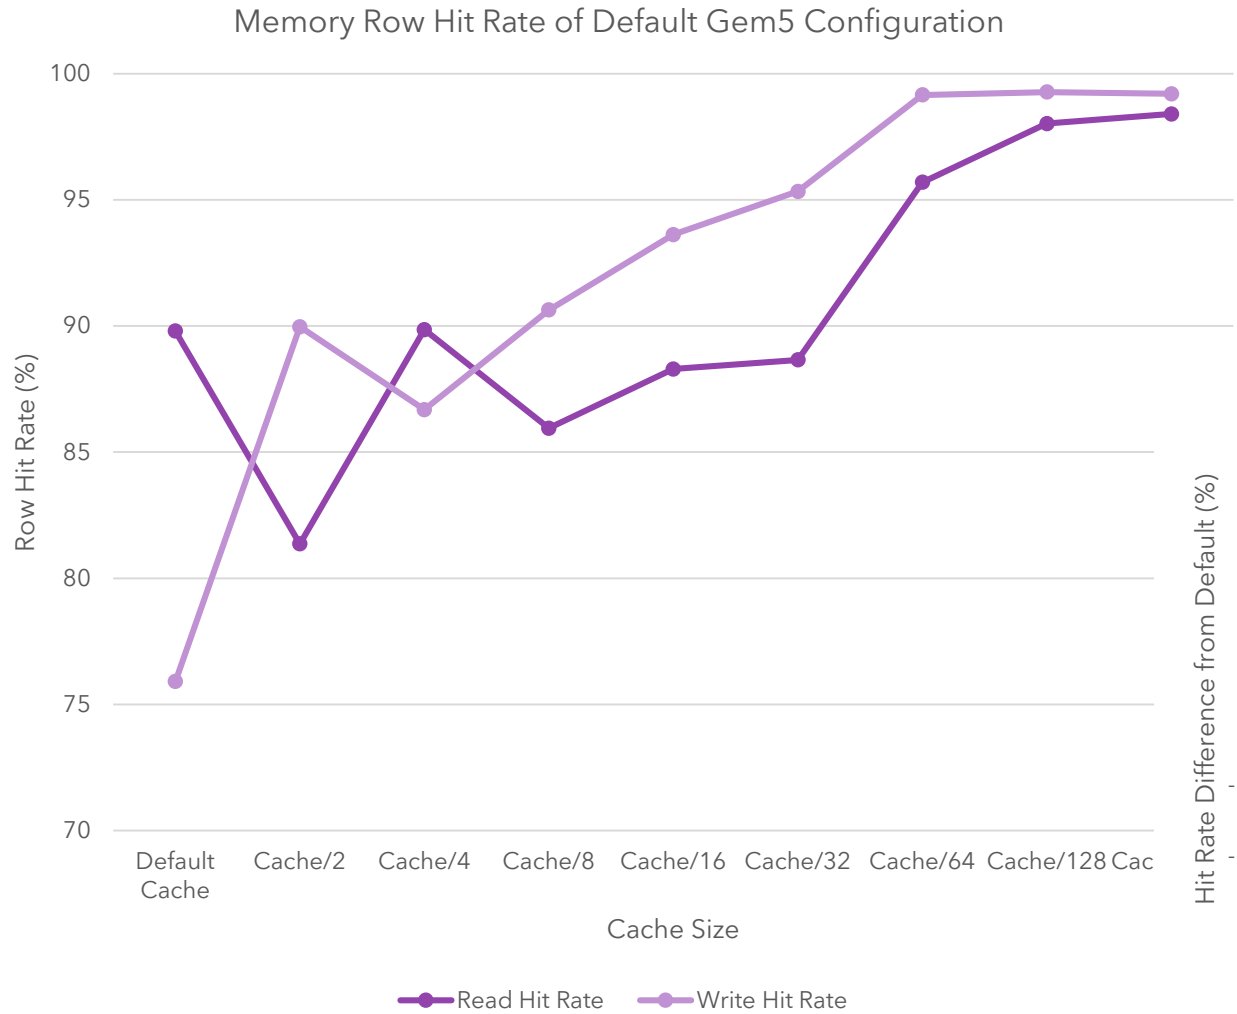


Activation Energy vs Max_retries




Average Read Latency VS Max Retries





- Baseline Gem5 configuration behaved strangely as cache size decreased.
- Provides some indication as to why ZORO didn't perform as well



- About 7% improvement overall

Future Work

- Modify ZORO to potentially improve performance with more complex row tracking
- Would have liked to have done more than a single benchmark
- Try different DRAM configurations, not just the single one we did
- Should also go up in cache size to determine if there are any changes
- We need to take into account CPU metrics instead of just the DRAM metrics
- Also would have liked to have a physical implementation of ZORO

Conclusions

- ZORO takes advantage of FR-FCFS by quickly scheduling transactions that are already in an open row
- ZORO could possibly save energy while having similar performance or even providing a small increase
- ZORO does not work well for lower cache sizes because it ends up being less efficient than simple scheduling

Related Work

- Hurr and C. Lin did some work trying to implement a history-based predictor as a scheduling mechanism for the memory controller

This differed from what we did because we aimed to change how the CPU schedules memory transactions

References

- [1] Hur and C. Lin, "Adaptive History-Based Memory Schedulers for Modern Processors," in *IEEE Micro*, vol. 26, no. 1, pp. 22-29, Jan.-Feb. 2006, doi: 10.1109/MM.2006.1.
- [2] J. Lowe-Power *et al.*, "The gem5 Simulator: Version 20.0+", *arXiv [cs.AR]*. 2020.
- [3] S. Li, Z. Yang, D. Reddy, A. Srivastava and B. Jacob, "DRAMsim3: a Cycle-accurate, Thermal-Capable DRAM Simulator," in *IEEE Computer Architecture Letter*


Questions?



*Where'd we get
the data?*

- Hit rate: # row hits / # commands
- Energy: Output from DRAMSim3
- Read latency: Output from DRAMSim3

Number of Read Commands vs Read Latency Category and Max Retries



0-500 500-1000 1000-1500 1500-2000 2000-2500 2500-3000 3000-3500 3500-4000 4000-4500