ECE 571 – Advanced Microprocessor-Based Design Lecture 25

Vince Weaver

https://web.eece.maine.edu/~vweaver

vincent.weaver@maine.edu

6 November 2024

Announcements

- HW#8 was posted
  - Please read paper for class Friday
  - The `perf mem record` is more informational rather than a full assignment
- Useful readings:
  - "A performance and power comparison of modern high-speed DRAM arch" from MEMSYS 2018
  - "DRAM Refresh Mechanisms, Penalties, and Trade-offs" Bhati, Chang, Chishti, Lu and Jacob. IEEE

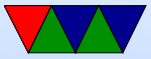

transactions on Computers, 2016.

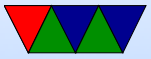

Some obscure/obsolete Memory Types

- RAMBUS RDRAM – narrow bus, fewer pins, could in theory drive faster. Almost like network packets. Only one byte at time, 9 pins?
- FB-DIMM – from intel. Mem controller chip on each dimm. High power. Requires heat sink? Point to point. If multiple DIMMs, have to be routed through each controller in a row.
- VCDRAM/ESDRAM – adds SRAM cache to the DRAM
- 1T-SRAM – DRAM in an SRAM-compatible package,

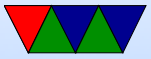

optimized for speed

- Hybrid Memory Cube – similar to High Bandwidth RAM (discontinued 2018)

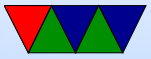

Memory Latencies, Labeling

- DDR400 = 3200MB/s max, so PC3200
- DDR2-800 = 6400MB/s max, so PC2-6400
- DDR2 5-5-5-15
  - C_L CAS latency
  - T_{RCD} row address to column address delay
  - T_{RP} row precharge time
  - T_{RAS} row active time
  - CMD (optional), command time
- DDR3 7-7-7-20 (DDR3-1066) and 8-8-8-24 (DDR3-1333).

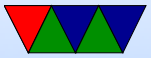

Memory Parameters

You might be able to set this in BIOS

- Burst mode – select row, column, then send consecutive addresses. Same initial latency to setup but lower average latency.
- CAS latency – how many cycles it takes to return data after CAS.
- Dual channel – two channels (two 64-bit channels to memory). Require having DIMMs in pairs

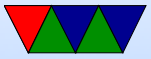

ECC Memory

- There's debate about how many errors can happen, anywhere from 10^{-10} error/bit*h (roughly one bit error per hour per gigabyte of memory) to 10^{-17} error/bit*h (roughly one bit error per millennium per gigabyte of memory
- Google did a study and they found more toward the high end
- Would you notice if you had a bit flipped?
- Scrubbing – only notice a flip once you read out a value

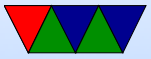

Registered Memory

- Registered vs Unregistered
- Registered has a buffer on board. More expensive but can have more DIMMs on a channel
- Registered may be slower (if it buffers for a cycle)
- Registered uses more power, and might be faster (even with the extra latency)
- RDIMM/UDIMM

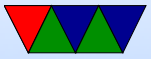

Bandwidth/Latency Issues

- Is RAM truly random access these days?
- No, burst speed fast, random speed not.
- Is that a problem? How is RAM accessed these days? Mostly filling cache lines?

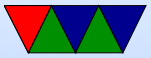

Memory Controller

- Do we even want full random access to memory?
  Should we just pass on CPU mem requests unchanged?
- What might have higher priority?
  Are some accesses more important? (regular vs prefetches)
- Why might re-ordering the accesses help performance (it's bad to switch back and forth between two pages)
- Can you improve mem performance with a better mem controller?

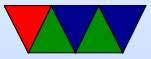

More on Refresh

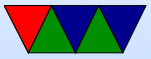

Reducing Refresh

- From *DRAM Refresh Mechanisms, Penalties, and Trade-Offs* by Bhati et al.
- Refresh hurts performance:
  - Memory controller stalls access to memory being refreshed
  - Refresh takes energy (read/write)

    On 32Gb device, up to 20% of energy consumption and 30% of performance

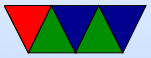

Async vs Sync Refresh

- Traditional refresh rates
  - Async Standard (15.6us)
  - Async Extended (125us)
  - SDRAM - depends on temperature, 7.8us normal temps (less than 85C) 3.9us above
- Traditional mechanism
  - Distributed, spread throughout the time
  - Burst, do it all at once (not SDRAM, just old ASYNC or LPDDR)

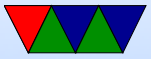

- Auto-refresh
  - Also CAS-Before RAS refresh
  - No need to send row, RAM has a counter and will walk the next row on each CBR command
  - Modern RAM might do multiple rows
- Hidden refresh – refresh the row you are reading? Not implemented SDRAM

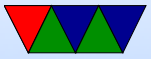

SDRAM Refresh

- Autorefresh (AR)
  - Device brought idle by precharging, then send AR (autorefresh)
  - Has a counter that keeps track of which row it is on, updates on each AR
  - The memory controller needs to send proper number of AR requests
  - LPDDR is a bit more complicated
  - Takes power, as all of SDRAM active while refreshing

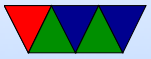

- Self-Refresh (SR)
  - Low-power mode
  - All external access turned off, clocks off, etc.
  - Has simple analog timer that generates clock for sending refresh pulses
  - Takes a few cycles to come out of SR mode

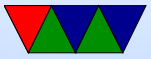

LPDDR Refresh features

- LPDDR has extra low-power features in SR mode
- Temperature compensated self-refresh (temp sensor)
- Partial-array self refresh (PASR), only refresh part of memory

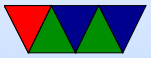

Refresh Timings

- Most SDRAM have 32 or 64ms retention time (tREFW)
- One AR command should issue in interval time (tREFI)
- A DDR3 with tREFI of 7.8us and tREFW of 64ms then 8192 refreshes
- Spec allows delaying refreshes if memory is busy

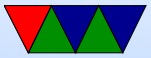

DRAM Retention Time

- Varies per-process, per chip
- Some chips over 1s, but have to handle worst-case scenario

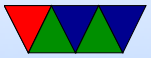

What can be done to improve refresh behavior

- Can you only refresh RAM being used? How do we know if values no longer important? free()? trim() command sort of like on flash drives?
- Probe chip at boot to see what actual retention time is, only refresh at that rate? Does chip behavior change while up?

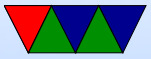

Refresh-related Security Issues – Rowhammer

- Been observed for years, adjacent rows discharging can affect nearby rows
- Particularly bad in DDR3 from 2012-2013
- Accessing same row over and over can make voltage fluctuations in nearby rows, causing faster leakage than normal

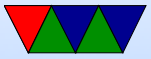

- Mitigations? Refresh more often? ECC? Refresh nearby lines if a lot of row hammering going on?
- Can cause exploit. Google NaCl disable "cflush" exploit (need to force access to row)
- Can also trigger just with lots of cache misses
- If you can flip bits of kernel/trusted pointers to point to something you control, then you win.

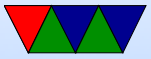

Advanced/Recent DRAM Developments

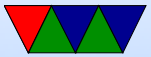

Cryogenic Memory

- Dip DIMMS in liquid nitrogen
- Low power? Faster? Interface with quantum circuits?