Final

What the exam looks like.

- Simple definitions, comparisons, advantages & disadvantages
  - what is it?
  - how does it work?
  - why have it?
  - pay particular attention to the terminology that was highlighted in the slides
- Apply the concepts and techniques you have learned to situations you have (hopefully) not seen before.

The goal is to test your knowledge of the material and how well you can apply it, **not how fast you can tell me what you know**.

The final will be a little longer than the midterm, with twice the test time.

Evan and Dustin will take the exam beforehand to make sure that they can finish it in a little **less than** half the time you will have.

Topics

Pipelining

- the basics
  - what is it
  - how it works
  - what hardware is involved
  - advantages/disadvantages
  - performance impact
- the 3 hazards (structural, data, control)
  - their causes
  - **all the solutions, their hardware requirements (if any), what cycles they save, what stall cycles they cannot eliminate**
    - forwarding
    - pipelined interlocks
    - pipeline flushing
  - special cases:
    - load-use hazard (delayed loads)
  - dependences vs. hazards
- branch prediction
- the R2000 pipeline structure

Topics

Advanced pipeline design

- superscalars & superpipelining
  - what are they
  - how they work
  - advantages/disadvantages
  - performance impact

Code scheduling

- instruction-level parallelism
- hiding operation latencies

Architectures & pipelines

Topics

Caches

- cache hierarchies
  - what problem they are solving
  - why they work
  - how they are typically configured & why
- locality of reference
  - 2 types
- caches
  - RAM components & their function
  - **accessing a cache**
  - **cache configuration**
  - **cache design trade-offs**
    - cache size
    - block size
    - associativity
    - memory update policy
    - block replacement policy
    - content
    - virtually/physically accessed
  - misses (the 3 C's)

Topics

Memory management

- relocation, base & bounds, fragmentation, overlays
- paging
  - advantages & disadvantages
  - virtual address space, physical address space
  - pages, page frames
  - **address translation**
  - page tables
    - content (PTEs)
    - **accessing**
    - **size**
  - handling page faults
- TLBs
  - why we have them
  - **how they're accessed**

Comparing caches & TLBs & paging

Topics

Performance

- pipelines
  - execution time
  - speedup
- memory systems
  - hit/miss rates
  - miss penalty
  - average (effective) memory access time for a single cache & a cache hierarchy
  - memory stall cycles in execution time

Topics

Exception

- how they’re handled
- hardware support
  - status register & vectored interrupts
  - effect on the pipeline