Midterm Review

1   Time and Location

The midterm will be given during normal class hours, 2:00p.m.-3:15p.m., on Monday October 15th, in the normal meeting place, EGR 0108.

2   Format

The midterm will be open-book, open-notes, closed everything else (you can bring any material that you have accumulated in class, but nothing else). Expect problems that will require you to analyze designs, compute performance, analyze the execution of code, etc. It is recommended that you bring a calculator.

3   Scope

The midterm will cover all material from the first day of class up and including lecture 13, given on October 10th. Essentially, this includes the introductory material in Chapter 1 of H&P, basic pipelining, introduction to ILP, superscalar processors, scoreboarding, Tomasulo's algorithm, and dynamic branch prediction. We will probably start caching techniques on October 10th, but this won't be emphasized on the exam.

4   Course Content

Here's a fairly comprehensive outline of the topics covered thus far. Disclaimer: this is not meant to be an absolutely water-tight complete list of topics. In other words, if there is a topic *not* present in this list, it may still show up on the midterm. However, it is a pretty good first-cut at what we have covered.

```
I.  Preliminaries
    A.  Architecture basics
        1.  Technology
        2.  Applications
        3.  Cost
        4.  Interface design
        5.  Performance evaluation
```

```
    B.  Performance basics
        1.  Amdahl's law
        2.  Performance equation
            T = I * CPI * tcycle


II.  Basic Pipelining Review
    A.  Implementing an ISA
        1.  Computer architecture building blocks
        2.  Pipeline implementation of MIPs
            a.  Stages (F,R,E,M,W)
            b.  Pipeline registers
            c.  Computing pipeline performance
    B.  Structural hazards
    C.  Data hazards
        1.  Types (RAW,WAW,WAR)
        2.  Pipeline stall
        3.  Bypass (forwarding)
    D.  Control hazards
        1.  Pipeline stall
        2.  Early branch test and target calculation
        3.  Static prediction
        4.  Delayed branch
    E.  Exceptions
        1.  Types (synchronous, asyncronous)
        2.  Restartability
        3.  Precise semantics
        4.  Problems with synchronous exceptions
            a.  Multiple exceptions
            b.  Out-of-order exceptions
        5.  Single commit point
    F.  Multi-cycle operations
        1.  Bandwidth, latency
        2.  Impact on pipeline
            a.  New structural hazards (competing for writeback port)
            b.  More forwarding paths
            c.  New data hazards that must stall
            d.  Complications for precise interrupts


III.  Instruction-Level Parallelism
    A.  ILP basics
        1.  Dependences
            a.  Data
            b.  Name
            c.  Control
        2.  Code scheduling
        3.  Memory disambiguation
        4.  Register renaming
        5.  Loop unrolling
    B.  In-Order Superscalar
        1.  Dynamic issue
        2.  Hazards (more structural hazards possible)
        3.  Issue rules
```

C. Dynamic instruction scheduling
        1. General concepts
            a. Out-of-order execution (/issue)
            b. Dynamic loop unrolling
            c. Instruction window size
        2. Scoreboarding
            a. Busy bits
            b. Scoreboard structure
            c. Checking for structural hazards
            d. Checking for WAW hazards
            e. Checking for WAR hazards
            f. Limitations
        3. Tomasulo's algorithm
            a. Ready bits
            b. Reservation stations
            c. Register renaming (register dataflow)
            d. Result bus
            e. Conflict queue (memory dataflow)
        4. Code scheduling examples using scoreboarding or Tomasulo
    D. Dynamic branch prediction
        1. Branch history table
        2. Multi-bit predictors
        3. Correlating predictors
        4. Branch target buffers

IV. Memory Systems
    A. Caches
        1. Principle of locality
        2. Cache organization
            a. Direct-mapped
            b. Set associative
            c. Fully associative
            d. Addressing the cache
        3. Cache Management
            a. Replacement policies
            b. Write-hit policies
            c. Write-miss policies
            d. Instructions vs. data