

EXAM

EECS 573: Microarchitecture

Instructor: Todd Austin

April 4, 2007

Name: _____

Score: _____

This is an open book, open note examination taken under the conditions of the University of Michigan honor code. You will have 90 minutes to complete the exam; keep careful track of your time to insure that you have the opportunity to answer each question. The exam totals to 55 points and will constitute 30% of your final grade in the course. Good luck!

Honor Pledge (please write it out and sign it):

2. Please answer the following questions regarding high-performance dynamic instruction scheduling:

a. [5 points] Draw a figure of a modern pipelined dynamic scheduler, detailing each stage of the scheduler and the replay loop. Describe why the replay loop is required, and what would be the cost of eliminating it.

b. [5 points] The select-free scheduler uses a predictor to reduce tag-matching requirements during wakeup. Describe this predictor and list the program and runtime characteristics that affect its accuracy. What would be the downside (to Brown's approach) if this predictor were eliminated?

- c. [5 points] Assume value prediction (predictors that predict the inputs to instructions) were to become prominent (and effective) speculation mechanisms in future machines, how might this affect the utility and design of instruction scheduling mechanisms?
- d. [5 points] What aspect of the non-stalling counterflow architecture makes it more scalable than a conventional microarchitecture? In the design, results must travel through at least $\frac{1}{2}$ of the result track entries before being written back to the ROB, why is this necessary? What characteristic of machine efficiency does counterflow adversely affect (be descriptive)?

3. Please answer the following questions regarding low-power microarchitecture design:

a. [5 points] Describe the three components of circuit power dissipation. For each power dissipation component, describe one optimization that can reduce the component. One of the optimizations that you suggest must be implemented in the compiler.

b. [5 points] Dynamic zero compression is a technique to reduce dynamic power dissipation in on-chip cache memory arrays. Describe how the technique saves dynamic power, and detail what are the critical design parameters for the technique and how they (qualitatively) affect performance, power dissipation, and area cost. Could this approach be adapted to reduce static (leakage) power? If so, how? If not, why not?

4. Please answer the following questions regarding circuit-sensitive design:

- a. [5 points] Detail (using a figure) why a direct mapped-cache has faster hit latency than a set-associative cache. Which type of cache would have better overall memory access latency for very small cache sizes? Which type of cache would have better overall memory access latency for very large caches sizes (e.g., multi-megabyte)? Explain.

- b. [5 points] How does the balanced tag cache mitigate the speed difference between the direct mapped cache and that set-associative cache? According to Peir's analyses, which cache has a faster access cycle? Which cache has better overall performance, and why?

5. [5 points] The quality of the research papers we studied this semester varied dramatically. Please give at least three pitfalls to avoid when performing and presenting technical research.