

This test is open book and open notes. Calculators are allowed.

HINT: Brief answers are good. Show your work for calculations. Sentence form answers should be short and to the point. Irrelevant responses may get negative scores.

Do any SIX of the seven questions below. If you write solutions or partial solutions to all seven, we will arbitrarily select six to grade unless you specifically identify which problem not to grade.

1. (18 points) This question is about the paper: “Exploiting Instruction Level Parallelism in Processors by Caching Scheduled Groups.”
 - (a) (5 points) Summarize the main idea of the paper in one or two sentences. In other words, how do the authors propose to get high performance?
 - (b) (5 points) Write two or three sentences to describe how register renaming is performed with DIF. Describe one advantage compared with the traditional (i.e. MIPS R10000 approach) (one sentence). Describe a situation where the traditional approach could work better (one sentence).
 - (c) (4 points) Give one reason from the paper that pure VLIW architectures perform poorly on commercial workloads. Does this consideration affect the IA-64?
 - (d) (4 points) The authors suggest that a DIF machine could achieve good performance without an L1 I-cache. How could this observation be used to simplify the design of the primary engine?

2. (18 points) This question is about the paper: “A 0.18 μ m CMOS IA-32 Processor With a 4-GHz Integer Execution Unit.”
 - (a) (6 points) Summarize the main idea of the paper in one or two sentences. In other words, how do the authors propose to get high performance?
 - (b) (3 points) Based on the data given in the paper for Intel processors, by what factor can clock frequency be increased when the number of stages in the pipeline is doubled?
 - (c) (3 points) A common instruction sequence is **add** \rightarrow **load** \rightarrow **test** \rightarrow **branch**, where the **add** operation is address calculation for the load, and the **test** operation is an add or subtract to compare two values. Write one or two sentences to explain how the result of the **add** is used by the **load**. How many clock cycles at what frequency elapse from starting the **add** until starting the **load**?
 - (d) (3 points) For the same instruction sequence as part c, write two or three sentences to explain how the result of the **load** is used by the **test**. How many clock cycles at what frequency elapse from starting the **load** until starting the **test** if the **load** hits in the L1 data cache? What happens if the **load** misses in the L1 data cache?
 - (e) (3 points) For the same instruction sequence as part c, write two or three sentences to explain how the result of the **test** is used by the **branch**. How many clock cycles at what frequency elapse from starting the **test** until starting the **branch**? How does this impact overall processor performance?

3. (18 points) This question compares the papers: “Exploiting Instruction Level Parallelism in Processors by Caching Scheduled Groups” and “A 0.18 μ m CMOS IA-32 Processor With a 4-GHz Integer Execution Unit.”
- (a) (10 points) What feature of the Pentium-4 most closely corresponds to a DIF cache? State two points of similarity between these two features, and state two differences. Your answer can be in point form; full sentences are not expected.
 - (b) (4 points) Compare how DIF and the Pentium-4 handle complex instructions. A two or three sentence answer is adequate.
 - (c) (4 points) Compare how DIF and the Pentium-4 provide precise exceptions. A two or three sentence answer is adequate.
4. (17 points)
- (a) (5 points) How does the complexity of a super-scalar processor increase with the issue width? Give an asymptotic form (e.g. $O(\log n)$, $O(n)$, $O(n!)$, etc.). State two aspects of super-scalar design that lead to this complexity.
 - (b) (3 points) How does the IA-64 architecture attempt to deal with the complexity from part (a)? Your answer should consist of one or two sentences.
 - (c) (3 points) How does the SMT architecture attempt to deal with the complexity from part (a)? Your answer should consist of one or two sentences.
 - (d) (3 points) How does the CMP architecture attempt to deal with the complexity from part (a)? Your answer should consist of one or two sentences.
 - (e) (3 points) How does the DIF architecture attempt to deal with the complexity from part (a)? Your answer should consist of one or two sentences.
5. (17 points) For each term below, give a one sentence definition.
- (a) (3 points) ICOUNT.
 - (b) (4 points) Predicated execution.
 - (c) (3 points) Victim caching.
 - (d) (4 points) WAR hazard.
 - (e) (3 points) Zoning.

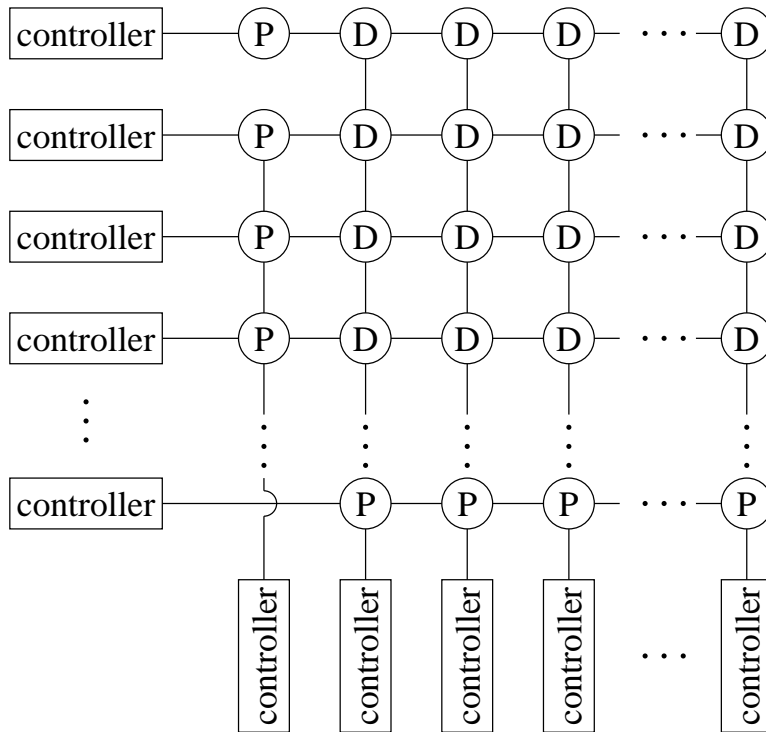


Figure 1: A Modified RAID Array

6. (17 points) This question pertains to the modified RAID array shown in figure 1. Each circle depicts a disk. Disks labeled with a “D” hold data, and those labeled with a “P” hold parity for their row or column. Each disk has interfaces to two controllers, one for its row, and one for its column.
- (3 points) Explain why this array can survive any two disks failing without loss of data or availability.
 - (3 points) Describe a three disk failure that results in loss of data.
 - (3 points) Explain why this array can survive any controller failing without loss of data or availability.
 - (8 points) Consider an array where each row has 20 data disks and one parity disk (except for the bottom row, which has 20 parity disks), and each column has 20 data disks and one parity disk (except for the left column, which has 20 parity disks). Assume that disks have a mean-time-to-failure of 500,000 hours, and that disk failures are Poisson distributed. For simplicity, assume that disks are the only things that ever fail in this array (e.g. controllers never fail). Assume that once per day, any faulty disks are replaced.

What is the mean-time-to-data-loss for this array?

You may make and state reasonable assumptions to further simplify the problem and still receive full credit as long as your final answer is within a factor of two of the exact answer.

7. (17 points) An argument for simple architectures such as CMP and DIF is that the simpler design should support a higher clock rate. On the other hand, processors such as the Pentium-4 show that a complicated superscalar can achieve high clock rates by making the pipeline deep enough. One consequence of deep pipelining is a higher branch mis-predict penalty. This problem explores the effect of branch mis-predicts.

(a) (3 points) Using the data from the paper “The Case for a Single Chip Multiprocessor,” what is the average instructions-per-cycle achieved by the two-issue superscalar and the six-issue superscalar for the nine benchmark programs that the authors considered?

(b) (12 points) The paper doesn’t report branch mis-predictions. For simplicity, assume that both processors achieved perfect branch prediction for the data reported in the paper. Assume that the 2-issue machine has a relatively short pipeline with a three cycle branch mis-predict penalty. Assume that the 6-issue machine has a long pipeline with a twenty cycle branch mis-predict penalty. Assume that 12% of instructions are branches or indirect jumps (i.e. require prediction). Assume that both processors operate at the same clock rate.

For what branch mis-predict rate do the two machines have the same performance?

(c) (2 points) When the mis-predict rate increases, does the ratio

$$\frac{\text{CMP performance}}{\text{super scalar performance}}$$

increase or decrease?