

- 4 (40 pts) A pipeline processor has $N = 20$ segments, and can throw a Type 1 exception in Segment #5, and can also throw a Type 2 exception in Segment #17. In both types of exceptions, the cost of flushing the pipeline is 1 cycle, and the cost of handling the exception is 1 cycle. Given these assumptions, please answer the following questions, and show all work:
- a) If the number of instructions in program P executed on this pipeline is given by $M = 8,998$, then how long will it take to execute P if the probability of Type 1 exceptions is 0.0012 and the probability of Type 2 exceptions is 0.0056?
- b) Suppose we optimize the pipeline to handle Type 1 exceptions in 1 cycle, but Type 2 exceptions are now handled in 3 cycles (still taking 1 cycle to flush the pipeline in either case). By what *percent* will the performance of this pipeline on program P (from part a), above) increase or decrease?
- 5 (30 pts) A program P has four types of instructions. On machine M, Type 1 has $IC = 28$ and $CPI = 2.3$, Type 2 has $IC = 37$ and $CPI = 1.8$, Type 3 has $IC = 72$ and $CPI = 2.4$, and Type 4 has $IC = 86$ and $CPI = 1.9$. Calculate the average CPI.

- 6 (30 pts) A pipeline processor has N segments and runs a program with $M = 1,298$ instructions, of which 22 percent are branches. Please answer the following questions, and show all work:
- Assuming that the branch prediction strategy is *branch not taken* and there is no branch delay slot, and there is a 50 percent probability that a branch will be taken, then what is the average CPI for this pipeline?
 - Assuming that the branch prediction strategy is *branch taken*, and there is a branch delay slot with a 48 percent probability that a branch will be taken, then what is the average CPI for this pipeline?
- 7 (30 pts) A cache memory is direct-mapped, and stores $N_D = 1$ Mbyte of data, with $L = 16$ words per block. Answer the following questions, and show all work:
- How many blocks are in the cache, and how many bits in the cache address if there is one valid bit per cache entry?
 - To which block, or range of blocks, does the byte address $A = 733,894$ map?
 - What is the total cache size, in bytes?

8 (30 pts) A cache memory is direct-mapped, and stores $N_D = 2$ Mbyte of data, with $L = 8$ words per block. Let the memory incur time delay $t_A = 1$ cycle to transmit the address, $t_M = 11$ cycles for each DRAM access, and $t_D = 1$ cycle to transmit each data word from memory. Answer the following questions, show all work:

a) Calculate the miss penalty C_{miss} , in *cycles*:

b) Using the cache described above, with a miss penalty of 23 cycles, assume that a program P has $IC = 1,415$ instructions, and 273 of these are loads (i.e., memory-to-register transfers) and 1.3 percent of these loads generate cache misses. Further assuming that the remaining instructions do not generate stalls, branch hazards, cache misses or any other type of exception, then calculate the CPI of a 7-segment pipeline running program P.

9 (10 pts **EXTRA CREDIT**) Give an example of the MIPS design rule *Simpler is Faster*, and how it applies to design of the MIPS ISA:

10 (20 pts **EXTRA CREDIT**) Let a CPU have an enhanced floating-point unit F that causes the CPU to run a program P twice as fast as before F was enhanced. If F is used by program P at least 60 percent of the time, then how fast does F have to be accelerated in order to achieve this 2X speedup? (Please show all work)

11 (15 pts **EXTRA CREDIT**) Compare and contrast (quantitatively) the effects of CPI three branch handling schemes that we discussed in class: (1) branch not taken, (2) branch taken, and (3) branch-taken with branch delay slot. Use equations and numerical examples for full credit.
