

Q2.1 [5] < 2.1> We wish to compare the performance of two different machines: M1 and M2. The following measurements have been made on these machines:

Program	Time on M1	Time on M2
1	10 seconds	5 seconds
2	3 seconds	4 seconds

Which machine is faster for each program and by how much?

Q2.2 [5] < 2.1> Consider the two machines and programs in Exercise 2.1. The following additional measurements were made:

Program	Instructions executed on M1	Instructions executed on M2
1	200 million	160 million

Find the instruction execution rate (instructions per second) for each machine when running program 1.

Q2.3 [5] < 2.2-2.3> If the clock rates of machines M1 and M2 in Exercise 2.1 are 200 MHz and 300 MHz respectively, find the clock cycles per instruction (CPI) for program 1 on both machines using the data in Exercise 2.1 and 2.2.

Q2.4 [5] < 2.2-2.3> Assuming the CPI for program 2 is the same as the CPI for program found in Exercise 2.3, find the instruction count for program 2 running on each machine using the exercise times from Exercise 2.1.

Q2.10 [5] < 2.2-2.3> Consider two different implementations, M1 and M2, of the same instruction set. There are four classes of instructions (A, B, C, and D) in the instruction set.

M1 has a clock rate of 500 MHz. The average number of cycles for each instruction class on M1 is as follows:

Class	CPI for this class
A	1
B	2
C	3
D	4

M2 has a clock rate of 750 MHz. The average number of cycles for each instruction class on M2 is as follows:

Class	CPI for this class
A	2
B	2
C	4
D	4

Assume that peak performance is defined as the fastest rate that a machine can execute an instruction sequence chosen to maximise that rate. What are the peak performances of M1 and M2 expressed as instructions per second?

Q2.11 [10] < 2.2-2.3> If the number of instructions executed in a certain program is divided equally among the classes of instructions in Exercise 2.10, how much faster is M2 than M1.

Q2.12 [5] < 2.2-2.3> Assuming the CPI values from exercise 2.10 and the instruction distribution from Exercise 2.11, at what clock rate would M1 have the same performance as the 750-MHz version of M2?

Q2.13 [10] < 2.2-2.3> Consider two different implementations, M1 and M2, of the same instruction set. There are three classes of instructions (A, B, and C) in the instruction set. M1 has a clock rate of 400 MHz, and M2 has a clock rate of 200 MHz. The average number of cycles per instruction (CPI) for each class of instruction on M1 and M2 is given in the following table:

Class	CPI on M1	CPI on M2	Instruction mix for C1	Instruction mix for C2	Instruction mix for C3
A	4	2	30%	30%	50%
B	6	4	50%	20%	30%
C	8	3	20%	50%	20%

The table also contains a summary of how three different compilers use the instruction set. C1 is a compiler produced by the makers of M1, C2 is a compiler produced by the makers of M2, and C3 is compiler produced by an independent compiler vendor. Assume that each compiler uses the same number of instructions for a given program but that the instruction mix is as described in the table.

- i. Using C1 on both M1 and M2, how much faster can the makers of M1 claim that M1 is compared with M2?
 - ii. Using C2 on both M1 and M2, how much faster can the makers of M2 claim that M2 is compared with M1?
 - iii. If you purchase M1 which of the three compilers would you choose?
 - iv. If you purchase M2 which of the three compilers would you choose?
-

Q2.15 [10] < 2.2, 2.3, 2.7> We are interested in two implementations of a machine, one with and one without special floating-point hardware. Consider a program, P, with the following mix of operations:

floating-point multiply	10%
floating-point add	15%
floating-point divide	5%
integer instruction	70%

Machine MFP (Machine with Floating Point) has floating-point hardware and can therefore implement the floating-point operations directly. It requires the following number of clock cycles for each instruction class.

floating-point multiply	6
floating-point add	4
floating-point divide	20
integer instructions	2

Machine MNFP (Machine with No Floating Point) has no floating-point hardware and so must emulate the floating-point operations using integer instructions. The integer instructions all take 2 clock cycles. The number of integer instructions needed to implement each of the floating-point operations is as follows:

floating-point multiply	30
floating point add	20
floating-point divide	50

Both machines have a clock rate of 1000 MHz. Find the native MIPS ratings for both machines.
