CS152 Section 11

Q1: Memory Consistency Models

Consider the following two threads executing on two different cores. Assume that memory locations A, B, and C are all initialized to zero.

```
        P1:                     P2:
        li  x1, 1               li  x1, 2
I1      lw  x2, A       J1      sw  x1, B
I2      sw  x1, C       J2      lw  x2, C
I3      lw  x3, B       J3      sw  x1, A
```

We are interested in the final values of P1.x2, P1.x3, and P2.x2.

Q1.1: Sequential Consistency
Give all possible sets of values of P1.x2, P1.x3, and P2.x2 under sequential consistency (SC).

Q1.2: Weak Versus Strong Memory Consistency Models
In general, what is the difference between a weak and a strong memory consistency model?

Q1.3: W→R Relaxation
Give all new possible sets of values if we relax Write → Read ordering constraints and the instruction orderings that caused them.

Q1.4: W→W Relaxation
Give all new possible sets of values if we relax Write → Write ordering constraints and the instruction orderings that caused them.

Q1.5: R→R and R→W Relaxation
Give all new possible sets of values if we relax Read → Read and Read → Write ordering constraints and the instruction orderings that caused them.

Q2: True / False

Indicate whether the following statements are true or false:

1. Sequential consistency is guaranteed if all processors have in-order pipelines.

2. A high-level language with a sequentially consistent memory model can be implemented on a ISA with a weaker memory model if fence instructions are provided.

3. Suppose an ISA specifies a non-multi-copy-atomic memory model, but a particular hardware implementation provides sequential consistency. Will software written for this ISA execute correctly on this machine?