Answers to Check Yourself

Chapter 1

§1.1, page 10: Discussion questions: many answers are acceptable.

§1.4, page 24: DRAM memory: volatile, short access time of 50 to 70 nanoseconds, and cost per GB is $5 to $10. Disk memory: nonvolatile, access times are 100,000 to 400,000 times slower than DRAM, and cost per GB is 100 times cheaper than DRAM. Flash memory: nonvolatile, access times are 100 to 1000 times slower than DRAM, and cost per GB is 7 to 10 times cheaper than DRAM.

§1.5, page 28: 1, 3, and 4 are valid reasons. Answer 5 can be generally true because high volume can make the extra investment to reduce die size by, say, 10% a good economic decision, but it doesn't have to be true.

§1.6, page 33: 1. a: both, b: latency, c: neither. 7 seconds.

§1.6, page 40: b.

§1.10, page 51: a. Computer A has the higher MIPS rating. b. Computer B is faster.

Chapter 2

§2.2, page 67: ARMv8, C, Java.

§2.3, page 73: 2) Very slow.

§2.4, page 81: 2) – 8ten

§2.5, page 89: 3) SUB X11, X10, X9

§2.6, page 92: Both. AND with a mask pattern of 1s will leaves 0s everywhere but the desired field. Shifting left by the correct amount removes the bits from the left of the field. Shifting right by the appropriate amount puts the field into the rightmost bits of the doubleword, with 0s in the rest of the doubleword. Note that AND leaves the field where it was originally, and the shift pair moves the field into the rightmost part of the doubleword.

§2.7, page 99: I. All are true. II. 1).

§2.8, page 110: Both are true.

§2.9, page 115: I. 1) and 2) II. 3).

§2.10, page 124: I. 4) ± 1024 K. II. 4) ± 128 M.

§2.11, page 128: Both are true.

§2.12, page 137: 4) Machine independence.

Chapter 3

§3.2, page 191: 2.

§3.5, page 229: 3.

Chapter 4

§4.1, page 260: 3 of 5: Control, Datapath, Memory. Input and Output are missing.

§4.2, page 263: false. Edge-triggered state elements make simultaneous reading and writing both possible and unambiguous.

§4.3, page 270: I. a. II. c.

§4.4, page 283: Yes, Branch and ALUOp0 are identical. In addition, you can use the flexibility of the don't care bits to combine other signals together. ALUSrc and MemtoReg can be made the same by setting the two don't care bits of MemtoReg to 1 and 0. Reg2Loc and RegWrite can be made to be inverses of one another by setting the don't care bit of Reg2Loc to 0. You don't need an inverter; simply use the other signal and flip the order of the inputs to the Reg2Loc multiplexor!

§4.5, page 296: 1. Stall due to a load-use data hazard of the `LDUR` result. II. Avoid stalling in the third instruction for the read-after-write data hazard on `X1` by forwarding the `ADD` result. 3. It need not stall, even without forwarding.

§4.6, page 309: Statements 2 and 4 are correct; the rest are incorrect.

§4.8, page 335: 1. Predict not taken. 2. Predict taken. 3. Dynamic prediction.

§4.9, page 342: The first instruction, since it is logically executed before the others.

§4.10, page 355: 1. Both. 2. Both. 3. Software. 4. Hardware. 5. Hardware. 6. Hardware. 7. Both. 8. Hardware. 9. Both.

§4.12, page 365: First two are false and the last two are true.

§4.13, page 4.13-4: Only statement #3 is completely accurate

§4.13, page 4.13-6: Statements #1 and #4 are true

Chapter 5

§5.1, page 391: 1 and 4. (3 is false because the cost of the memory hierarchy varies per computer, but in 2016 the highest cost is usually the DRAM.)

§5.3, page 412: 1 and 4: A lower miss penalty can enable smaller blocks, since you don't have that much latency to amortize, yet higher memory bandwidth usually leads to larger blocks, since the miss penalty is only slightly larger.

§5.4, page 431: 1.

§5.7, page 465: 1-a, 2-c, 3-b, 4-d.

§5.8, page 472: 2. (Both large block sizes and prefetching may reduce compulsory misses, so 1 is false.)

§5.11, page 5.11-8: All are true.

Chapter 6

§6.1, page 518: False. Task-level parallelism can help sequential applications and sequential applications can be made to run on parallel hardware, although it is more challenging.

§6.2, page 523: False. *Weak* scaling can compensate for a serial portion of the program that would otherwise limit scalability, but not so for strong scaling.

§6.3, page 528: True, but they are missing useful vector features like gather-scatter and vector length registers that improve the efficiency of vector architectures. (As an elaboration in this section mentions, the AVX2 SIMD extensions offers indexed loads via a gather operation but *not* scatter for indexed stores. The Haswell generation x86 microprocessor is the first to support AVX2.)

§6.4, page 533: 1. True. 2. True.

§6.5, page 537: False. Since the shared address is a *physical* address, multiple tasks each in their own *virtual* address spaces can run well on a shared memory multiprocessor.

§6.6, page 545: False. Graphics DRAM chips are prized for their higher bandwidth.

§6.7, page 550: 1. False. Sending and receiving a message is an implicit synchronization, as well as a way to share data. 2. True.

§6.8, page 552: True.

§6.10, page 564: True. We likely need innovation at all levels of the hardware and software stack for parallel computing to succeed.

§6.9, page 6.9-10: 1. pollingwith loads and stores. 2. interrupts with DMA.