CS 201 Section 1 Homework 4 Suggested Solutions

1. (5.14) Assuming the ordinary precedence of operators (`*` and `/` before `+` and `-`), the parse tree for `(2*3+4) - (4/2+1)` looks like this:

```
           -
          / \
         /   \
        /     \
       +       +
      / \     / \
     /   4   /   1
    *       /
   / \     / \
  2   3   4   2
```

The corresponding RPN is given by taking the nodes of the tree is post-order:

```
2 3 * 4 + 4 2 / 1 + -
```

Assuming all the numeric literals are meant to represent integers, the corresponding JVM code is:

```
ICONST_2
ICONST_3
IMUL
ICONST_4
IADD
ICONST_4
ICONST_2
IDIV
ICONST_1
IADD
ISUB
```

2. (5.32) The disk can read the contents of one track, i.e., $64 \times 512 = 32768$ bytes, on each complete rotation. Hence the controller generates data at a rate of 32768 bytes per 16 msec = 2048 bytes/msec. Data is sent in distinct transfers of 16 bits each, so the controller will steal 1024 bus cycles per msec. Since each cycle takes 500 nsec, the CPU will be blocked from executing for 512000 nsec per msec; i.e., it will be slowed down by about 51.2% (The fact that each CPU instruction takes an average of 2 bus cycles is irrelevant.)

3. (5.33) Interrupt service routines are called asynchronously under the stimulus of external events; several of them may want to be active at once. Priority levels are a mechanism for controlling the order in which each ISR gets to use the CPU. Ordinary subroutines are

called synchronously under program control; the order in which subroutines use the CPU is completely determined by the program's calling structure, so there's no need for priorities.

4. (5.36) The limiting factor is the time taken to process each packet interrupt. So the bridge can transfer one packet per 250 μsec, i.e., 4000 packets per second, 2000 from each network. At 1K bytes per packet, this limits each network to 2Mbytes/sec.

5. (2.9.) Each pixel requires $3 \times 6 = 18$ bits, so a single visual frame contains 1.8×10^7 bits. At 10 frames per second, the data rate is 1.8×10^8 bits/second. The brain doesn't process visual data that fast, however, according to Tanenbaum.

6. (2.28.) The display must draw $1024 \times 768 \times 75$ pixels per second, so each pixel pulse is 1.69×10^{-8} seconds.

7. (2.29.) (See p. 97, para. 4.) The usual way to handle this is to have a bank of 256 24-bit mapping registers in the controller hardware. Whenever a byte is fetched from video RAM, the 8-bit number is used as an index into the mapping register. The register selected delivers the 24 bits to drive the display (8 bits per color). Thus indeed 2^{24} colors are available, but at any instant, only 256 of them can be displayed. Changing the subset of displayable colors is done by reloading the mapping register.

8. (2.30.) There are 4000 characters per page. Each character covers (on average) $0.25 \times (2 \times 10^{-3}\text{m})^2 = 10^{-6}\text{m}^2$ and is 2.5×10^{-5}m thick, giving a volume of $2.5 \times 10^{-11}\text{m}^3$ of toner per character, hence 10^{-7}m^3 of toner per page. The capacity of the toner cartridge is $(25 \times 10^{-2}\text{m}) \times (8 \times 10^{-2}\text{m}) \times (2 \times 10^{-2}\text{m}) = 4 \times 10^{-4}\text{m}^3$. So the cartridge is good for $4 \times 10^{-4}\text{m}^3/10^{-7}\text{m}^3 = 4000$ pages.