HW9 Solutions

1. Memory Mapping

1.1

3 entries. The page size is 4096, so map1 consumes 2 entries, and map2 consumes one entry.

1.2

2 pages. The 3 virtual pages map to two physical pages in the OS buffer cache.

2. Thread Switching

```
void swtch(struct thread *t1, struct thread *t2) {
    // On entry this function is run by thread t1.

    push_register(RBP);
    push_register(R0);
    push_register(R1);
    push_register(R2);
    push_register(R3);

    t1->stack = read_register(RSP);
    write_register(RSP, t2->stack);

    pop_register(R3);
    pop_register(R2);
    pop_register(R1);
    pop_register(R0);
    pop_register(RBP);

    return; // The function should return to the
            // point where thread t2 called swtch.
}
```

3. Interrupts

3.1

200 words/min × 1 min/60 seconds × 7 letters/word × 1 interrupt/character = 23 interrupts/second.

3.2

It should use interrupts. Your interrupts cost the computer 23 microseconds out of every second; this is a trivial fraction (23 parts in one million). If the computer used polling, you (the human) would notice the lag and get annoyed.

4. Disk Calculations

4.1

20 GB/disk. 10 platters/disk × 4096 tracks/platter × 1024 sectors/track × 512 bytes/sector = 10×4×1024×1024×512 = 40×1MB×512 = 40×.5GB = 20 GB.

4.2

100MB/second. First note that 12,000 RPM = 200 rotations per second (or one rotation per 5 ms). In a single rotation, we can read an entire track. A track consists of 512 bytes/sector × 1024 sectors = 0.5 MB. So the sequential transfer bandwidth is 200 rotations/second × 0.5 MB/rotation = 100 MB/second. Because the track-to-track seek time and the I/O bus overhead are both modeled as negligible, 100 MB/second is our answer.

4.3

Roughly 30 KB/second. First note that in one read, we get 512 bytes. What is the time to issue this read? The disk incurs seek delay and rotational delay. The average seek latency is 15 ms. After the disk head reaches the desired track, the disk has to wait until the desired sector rotates under the disk head. Since the sector could be anywhere on the track, ranging from right under the head to the most pessimal position, the average rotational delay is 2.5 ms (half of the 5 ms per rotation). So the total delay on average is 17.5 ms. Our total effective bandwidth, then, is 512 bytes / 17.5 ms ≈ 525 / 17.5 bytes/ms = 30,000 bytes/second.

5. Disk Scheduling

Shortest-seek-time-first

20 - 22 - 10 - 6 - 2 - 38 - 40

So the time is (2 + 12 + 4 + 4 + 36 + 2) × 6 ms = 360 ms

LOOK

If travelling up first:
20 - 22 - 38 - 40 - 10 - 6 - 2

So the seek time is (2 + 16 + 2 + 30 + 4 + 4) × 6 ms = 348 ms

Or if it travels down first:
20 - 10 - 6 - 2 - 22 - 38 - 40

(10 + 4 + 4 + 20 + 16 + 2) × 6 ms = 336 ms