Midterm

LBSC 690
October 29, 2012

Name: _____
by writing my name I swear by the honor code

Read all of the following information before starting the exam:

- Show all work, clearly and in order, if you want to get full credit. I reserve the right to take off points if I cannot see how you arrived at your answer (even if your final answer is correct).
- You have two and a half hours to complete this exam (but you shouldn't need it).
- This exam is open book and open notes, but you cannot use a computer. You may use a calculator, but you shouldn't need to.
- Justify your answers algebraically whenever possible to ensure full credit. Be sure to have units for all answers that call for them.
- Circle or otherwise indicate your final answers.
- Please keep your written answers brief; be clear and to the point. I will take points off for incorrect or irrelevant statements.
- This test has six problems, but you only need to complete **four** questions. If you answer more than four questions, I will grade the first four completed questions. (If you change your mind on which questions you want graded, be sure to clearly identify which question shouldn't be graded.)
- Each question is worth ten points.
- Good luck!

1. You're in charge of providing students with tablets filled with 750 eBooks. The tablets don't have much memory—only 200MB. You thought that you were on the ball, double checking that the books that would go on the device would actually fit. While the files you checked do actually fit (they only take up 150MB), you checked plain ASCII files (no markup, just paragraphs separated by line breaks). What's actually going on the devices are EPUB files. Just to refresh your memory, EPUB books have:

   - An XHTML file for every chapter (or whatever the atomic unit of the book is)
   - Two XML files (a NCX file and an OPF file) with metadata and a linear order of the book

   The EPUB vendor hasn't given you the books yet, but they assure you that all of the XML files with the order and metadata are under 10KB per book. You're still worried that the books won't fit; let's see if they do.[1]

   (a) (3 points) How much in total do the extra XML metadata files take up for all the books in MB?

   ANSWER: $\frac{10\text{KB}}{\text{book}} \cdot \frac{1\text{MB}}{1\text{KB}} \cdot 750\text{books} = 7.5\text{MB}$

   (b) (5 points) How much more would an XHTML version of the book take up than the ASCII version? Make reasonable assumptions and explain why you made them. Give your answer as a conversion ratio (XHTML over ASCII). [Hint: this should be a unitless number greater than one (XHTML takes up more space) but probably less than two (XHTML doesn't double the size)].

   ANSWER: We need an XHTML page for every chapter. That XHTML will have a header and additional markup in the text. Let x be the number of chapters, let H be the size of the header, let y be the number of paragraphs per chapter, and let P be the amount of markup added for each paragraph. A reasonable assumption would be $H = .1K$ and $P = .01KB$ for books with $x = 100$ chapters and $y = 1000$ paragraphs per chapter (erring on the side of caution), that gives a total of 20KB per book. Given that the average book is about 200KB, that's a ratio of **1.1**.

   (c) (2 points) Given your above estimate, would the books fit?

   ANSWER: Given the answers to the previous parts, that's an additional 7.5 MB and 15MB, which is still less than 200 MB.

[1] EPUB files are typically compressed. To make life easier on you, ignore all compression. If you don't know what that is, that's totally fine (we haven't covered it yet).

2. Someone has been downloading big files from your library's web server, slowing down your Internet connection. To fix the problem, you look at the server's logs, but it's storing the IP address as a hexadecimal number. For instance, it would store 131.215.239.141 as 83D7EF8D.

   The computer that's doing the downloaded is recorded as 12091645; how would you write that in the conventional way of writing IP4 addresses? Make sure to explain how you got your answer.

   ANSWER: Every pair of hex digits corresponds to an octet.

   12 $1 \cdot 16 + 2 = 18$

   09 $0 \cdot 16 + 9 = 9$

   16 $1 \cdot 16 + 6 = 22$

   45 $4 \cdot 16 + 2 = 170$

   So the IP address is **18.9.22.170**.

3. Consider the vertical scroll bar of a web browser or word processor. For each of the interface principles below, give one way that the scrollbar uses the principle for effective design.

   (a) (2.5 points) Affordance

   ANSWER: Affordance means that the appearance of the interface should make the element's function obvious. Buttons are for pushing and live on either end of the scroll bar.

   (b) (2.5 points) Constraints

   ANSWER: The scroll bar can only move up and down, not left or right. When you've reached the end of the document, you can't move any further.

   (c) (2.5 points) Mapping

   ANSWER: The relative position of the scroll par matches the position of the document. As you move through the document, the bar moves closer the bottom of the scroll bar.

   (d) (2.5 points) Causality

   ANSWER: Any motion is the scroll bar, it is immediately reflected in the document.

4. In previous offerings of this class, I was able to talk about RAM and hard drives as two different things. Now things have gotten all confusing with these newfangled solid state drives (SSD).

   (a) (3 points) Here's a picture of one example of a SSD. I've circled three parts of it. What are they and what do they do? (Think about other pieces of the computer that we passed around in class.)

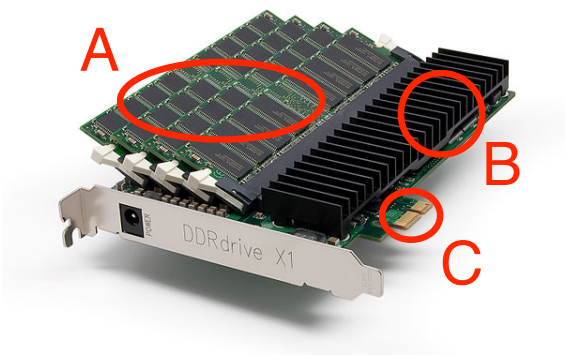

   ANSWER: There are RAM chips (A) attached to the main board. The heatsink (B) radiates off excess heat from the main board, which connects to the rest of the computer through the connector (C).

   (b) (3 points) This particular SSD needs to have a battery (not visible). Why is that? Do magnetic hard drives have a battery? Does RAM?

   ANSWER: Commodity RAM is volatile, which is why this SSD needs a battery to hold data long term when the computer is turned off. Magnetic hard drives do not have this, as they are not volatile. RAM typically does not need a battery because it's designed to forget its contents when the computer is powered down.

   (c) (4 points) I have a task that requires reading 10^9 sectors from my drive in a random order. How long would it take to do that with a latency of 0.05ms?

ANSWER: $10^9 \text{sectors} \cdot \frac{0.05\text{ms}}{sector} = 5 \cdot 10^9 \cdot 10^{-5} s = 50000s = 14 hours$

5. You have files laid out as follows:

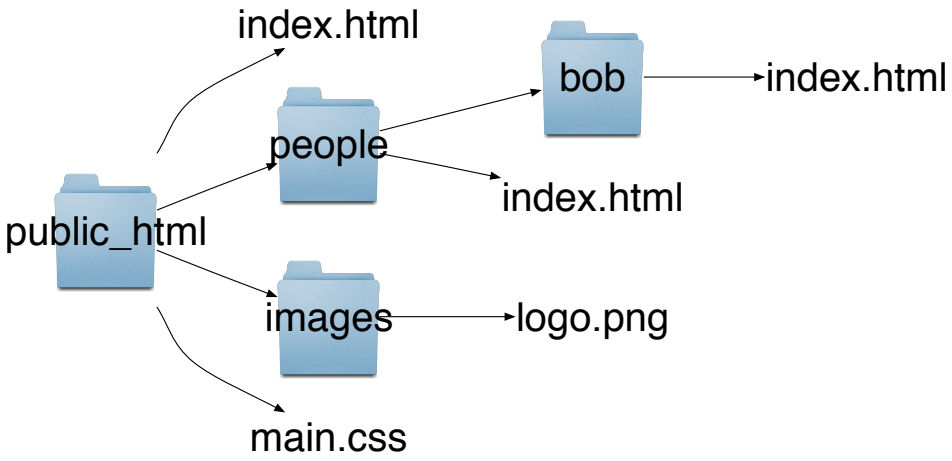

(a) (2.5 points) Every page includes the file main.css; what must be in that CSS file so that links: are red, aren't underlined by default, but are underlined when the user hovers over them?

```
a {
  text-decoration: none;
  color:#ff0000;
}
a:hover {
  text-decoration: underline;
}
```

(b) (2.5 points) What HTML would you need in the file index.html to use the CSS file?

```
<link rel="stylesheet" type="text/css" href="main.css">
```

(c) (2.5 points) What HTML would you need to display the image logo.png in the file index.html in the root directory?

```
<img src="images/logo.png" />
```

(d) (2.5 points) What HTML would you need to display the image logo.png in the file index.html in Bob's directory **and** make it link to the file index.html in the root directory?

```
<a href="../../index.html"><img src="../../images/logo.png" /></a>
```

6. One of the most contentious recent issues in Internet policy is *net neutrality*, which basically says that folks on the Internet should treat all packets as equal. For example, Comcast might have a streaming movie service that competes with Netflix. It might "slow down" Netflix's packets or speed up its own packets to make its offering more attractive.

   (a) (4 points) How does Comcast have the ability to do this? Explain in detail what aspects of IP packets allows them to do this (Hint: Netflix uses port 80 for its video—the same as for HTTP, so it's not as simple as just looking at the port number).

   ANSWER: Every IP packet has both a "from" field and a "to" field that gives the IP address of the server and and the requester. If Comcast has a list of a bunch of Netflix servers, they could do something malicious like holding Netflix packets with a delay or just give priority to their own Netflix-like services.

   (b) (3 points) Let's say that Comcast throttles Netflix to 1.0 Mbs. What is the largest movie file size that Netflix can stream in realtime? (As usual, make reasonable assumptions.)

   ANSWER: In two hours (a reasonable movie length), you'd get:
   $1\frac{Mb}{s}\frac{2\cdot 60\cdot 60s}{1}\cdot\frac{1\text{MB}}{8\text{Mb}} = 900\text{MB}$

   (c) (3 points) What bandwidth would Comcast need to stream DVD-quality movies?

   ANSWER: A DVD hold about 4GB, and if we assume a movie lasts two hours
   $\frac{4\text{GB}}{2\text{hours}} = \frac{32\cdot 10^3 Mb}{2\cdot 60\cdot 60s} = 4.4\text{Mbs}$