Answers to questions that students asked during the weekend before Exam 1

Q: Regarding the Question and Answer examples you gave us … are these a broad selection from which you will be pulling exam questions, or are they "model" questions from which the real questions might differ considerably? Basically, if I can answer all of the questions on my own, am I all set, or do I need to also review details for other potential questions?

A: They're just models. You need to reviews details of slides and required readings.

Q: In the lectures 6 and 7 (September 17 and 19, 2013), slide 19, what does "implementation-centric" mean in the Internet architecture?

A: The phrase "implementation-centric" is used in order to contrast the "real internet architecture" of slide 19 with the abstract OSI architecture of slide 18. In the OSI architecture, the physical, data-link, network, and transport layers are strictly defined in terms of the functionality, semantics, and data formats of the *layers*. In the real Internet, layers are more loosely defined (in terms of the *protocols* that are actually implemented and used), and the architecture is shaped like an hourglass, rather than a tower.

Q: In interdomain-routing slides 5 to 14 (September 24, 2013), what does the red arrow represent in the BGP convergence/oscillation scenario? And what's the difference between converge and oscillation in the same triangle topology of three nodes?

A: A red arrow from a source node to a neighbor indicates that the source node has chosen that neighbor as the first hop on its route to destination *d*. These slides present three scenarios that might unfold as these three nodes execute BGP. All three scenarios start with node *d*'s announcing itself to its neighbors 1 and 2; once they receive these announcements, source nodes 1 and 2 both know that they have the option of choosing direct, one-hop routes to *d*.

In the first scenario (slides 5-7), the second thing that happens (after *d*'s announcing itself to 1 and 2) is that node 2 selects its preferred route 2 → *d* and announces that route to node 1, its neighbor. Then node 1 selects its preferred route 1 → 2 → *d* (i.e., selects node 2 as the first hop on its route to *d*). At this point, BGP has *converged* on a confluent tree of routes to *d*.

In the second scenario (slides 8-11), the second thing that happens (after *d*'s announcing itself to 1 and 2) is that nodes 1 and 2 *simultaneously* select the one-hop routes 1 → *d* and 2→ *d* and announce these routes to each other. Although these are not the preferred routes of source nodes 1 and 2, they are the only routes available at this point, and both nodes would rather use them than not be able to reach *d* at all. The third thing that happens in this scenario is that both 1 and 2 *simultaneously* withdraw their one-hop routes, choose their preferred routes 1 → 2 → *d* and

2→ 1 → *d* and announce them. Right after that, both nodes realize that they cannot use their currently chosen routes, because doing so would result in a loop through 1 and 2 rather than a confluent tree to *d*; so they *simultaneously* withdraw their two-hop routes and re-announce the one-hop routes. As long as they keep making these decisions and announcements simultaneously, BGP may oscillate forever between {1 → 2 → *d*, 2 → 1 → *d*} and {1 → *d*, 2 → *d*}.

In the third scenario (slides 12-14), we have the same preferences as in the second but different timing. The second thing that happens (after *d*'s announcing itself to 1 and 2) is that node 2 selects 2 → *d* and announces it to node 1. The third thing that happens is that node 1 selects its preferred route 1 → 2 → *d*. Node 2 is never offered the option of routing through node 1, and so BGP converges on a confluent tree of routes to *d* (the same tree as in the first scenario).

Q: What are the answers to the questions on slide 25 of lecture 2 (September 3, 2013)?

A: These are, to some extent, "discussion" questions, but here's my stab at answering them.

The DMCA was supposed to be an *addition* to US copyright law as it existed in 1998; that is, all of the laws that were in effect before the DMCA passed are still in effect (unless they've since been repealed or amended).

In particular, fair use, which was allowed under US copyright law before the DMCA, is still allowed under US copyright law. Because fair-use advocates pointed out during the congressional debate over DMCA that copyright owners might use "effective technological-protection measures" to prevent fair use, certain exceptions to the DMCA's prohibition on circumvention of such measures were carved out – see slide 23 of lecture 2. Both the specific list of exceptions and the general approach (legislating a broad prohibition and then specifying an *ad hoc* set of exceptions) remain highly controversial, and fair-use advocates would probably say that the DMCA has in fact inhibited fair use to some extent during the 15 years that it has been in effect.

Yes, one can violate the DMCA without actually infringing any copyrights (by circumventing a technological-protection measure but not infringing the material that it is protecting).

Q: Do we really need to memorize some bullet points in exact details? For example, Internet protocol-design philosophy, fair-information principles, *etc.*?

A: You *do* need to know the items on these lists. However, you should strive to understand the bullet points, not just to memorize the words on the slides. If an exam question asks you to provide one of the lists we went over in class (or part of such a list), you will get credit for explaining the concepts in the list correctly even if you don't use exactly the same words that are on the slides.