**Answers to Comprehensive Exam: Networks (60 points) Autumn 2002**

1. (15 points total) *TCP transport protocol* TCP has a 32-bit sequence number field and a 16-bit window

   (a) ( 5 points) Give a quantitative argument, as the designers of TCP might have done in 1980, that 32-bits is a good choice of size for the sequence number field.

   *Answer:* The key issue is being able to detect old delayed segments showing up. The 32 bits allows 4 gigabytes to be transmitted before wrap, which is 8*4G/10M = 3200 seconds or almost one hour at 10 Mbps, which was high-speed back in those days. And, a TCP segment is not expected to live in the next more than a few seconds so this should catch old delayed segment even just using half the range.

   (b) (5 Points) In the same vein, why is 16 bits a reasonable choice for the window parameter.

   *Answer:* This allows up to 64 kilobytes to be outstanding, while on a LAN with 10 Mbps data rate, this allows a RTT of up to about 50 milliseconds, whereas the actual RTT is microseconds. WAN data rates tend to be much lower even though the RTT is longer.

   (c) (5 points) Quantify the limitations and concerns that these sizes raise based on how the technology has changed since 1980, i.e. lower cost memory, faster processors, optical links, etc.

   *Answer:* At 10 Gbps, you can wrap in 3.2 seconds, and usually one provides half the range for new packets and half for old, so you are in trouble if segments live from more than 1.6 seconds. With WANs, this is getting rather close to expected max. packet lifetimes, which can be hundreds of milliseconds. Similarly, for the window parameter, multi-gigabit links are arising that have 100 Ms RTT, so we cannot operate full-speed because we dont have the delay-bandwidth product.

2. (15 points total) *Virtual Circuits vs. Datagrams* Alexander Graham Bell returns from the dead to set us straight on (virtual) circuit switching, given he considers we went off the rails in using datagrams in the Internet. Give your cogent response to each of the points that this old cadaver makes:

   (a) ( 5 points) "Hot dang, this datagram stuff requires huge packet headers that incur excessive bandwidth overhead. Circuit switching allows you to use itsty bitsy tiny little circuit IDs."

   *Answer:* Large packets, which dominate the traffic, are 1000-1500 bytes and the headers are only 56 bytes or so, so percentage overhead is not significant. Even with average packet size of 300 bytes means the overhead is 20 percent. Moreover, lots of bandwidth is wasted with virtual circuit setup, and connections/flows in the Internet tend to be very short on average.

   (b) (5 Points) "This new fangled datagram nonsense cost much more to do packet classification and forwarding decisions in each switch because of the large clumsy headers. With circuit switching, the circuit id can be an index into a table that directly indicates the next hop."

*Answer:* Typically, only a small port of the header needs to be looked up for a forwarding decision, such as the 32-bite IPv4 address in a router, so the size is not much larger than a VCid. For higher-level classification, it would be more expensive to have a network-layer connection per higher-level connection (TCP) because they are so short-lived, and connections are short-lived — see above.

(c) (5 points) "This datagram nonsense does not allow you to reserve bandwidth so you cannot know whether your packets are going to get through. Good grief, the only guarantees are death and taxes, and I'm already dead."

*Answer:* There are no guarantees in either case, because a router can fail, At least with datagrams, your packets might get through by routing around the failure, rather than having to re-setup the whole virtual circuit. Moreover, you can do RVSP on top of datagrams if you insist on reservations. (Thanks for the telephone, in any case.)

3. (15 points total) *Ethernet*

(a) ( 5 points) Describe the Ethernet media access control (MAC) protocol, CSMA-CD, comparing it to the Aloha network and p-persistent protocols.

*Answer:* Listen before transmit, transmit if idle or wait for idle. Listen while transmitting and abort if collision and retry as above. Aloha does not listen for collisions. A p-persistent protocol only transmits with probability p if the line is idle, ie. Ethernet is 1-persistent.

(b) ( 5 points) Original Ethernet was at 10 Mbps with limitations on cable length and packet size. Now, IEEE is busy standardizing 10 Gbps Ethernet. Describe how these parameters need to change, if at all, and why if you insist on using the MAC protocol above at this higher speed.

*Answer:* You need to either decrease the max. cable length by a factor of 1000 or else increase the min packet size by a factor of a 1000 (or some combination of the two i.e. 100/10) so that transmitter is still listening/transmitting when there can be a collision. This is not practical to do. 64Kbytes min packet size!

(c) ( 5 points) Peterson and Davie say: "it might seem that a wireless protocol would follow the exactly the same algorithm as the Ethernet" as a lead-in to why not. Describe why not and what 802.11 does about it.

*Answer:* In the wireless environment, nodes cannot necessary hear the transmission of every other node that might interfere with its transmission, unlike Ethernet. The 802.11 solution is to request-to-send to the destination before sending, and getting clear to send back from destination, with the hope that others close by will see this and avoid colliding.

4. (15 points total) *End to end.*

(a) ( 6 points) Describe the "end-to-end" argument in networking, illustrating how a file transfer program should behave if it was truly "end-to-end".

*Answer:* End-to-end semantics can only be ensured by end-to-end mechanisms and checks because intermediate/lower-level mechanisms can fail in ways that only end-to-end mechanisms can detect. Lower-level mechanisms are at best, optimizations. An end-to-end file transfer ensures that a file was accurately transferred to a remote disk by transmitting a file-level checksum on the file to the receiver and having the receiver read back the file, recompute the checksum and verify against the supplied one. This catches any intermediate failures (except low prob. checksum fooling).

(b) (6 Points) Osama Bin Laden, attempting to point out yet another hypocrisy of the West, argues that if we really believed in end-to-end, we would not mess around with Ethernet CRCs, IP checksums, TCP checksums and TCP retransmissions. Describe why Bin Laden's wrong in this particular instance.

*Answer:* The TCP checksum and retransmission mechanism is an optimization over having to retransmit larger units, such as the whole file in a file transfer case, when a packet is dropped or corrupted. The Ethernet CRC helps further to detect corrupted packets, given that the TCP checksum is not very strong, and Ethernet CRC is normally computed in hardware. IP checksums might be argued helpful to avoid forwarding packets whose headers have been trashed, another optimization.

(c) (3 Points) Give an example of how Internet protocols are not completely consistent with the end-to-end argument.

*Answer:* Standard FTP does not include an end-to-end file-level checksum. Also, telnet has no application-level check other than the user noticing missing characters.

*The End*