Solutions

Comprehensive Exam: Networks (60 points) **Autumn 1999**

1. (15 points total) *End to end.*

   (a) (5 points) Describe the "end-to-end" argument in networking, given one good example illustrating it.

   Ans: Function needs to be provided at the endpoints, actually the application endpoint, to ensure application semantics of this functionality. Mechanisms in the "middle" (of the network) can only optimize. For example, a reliable end-to-end file transfer requires a file-level check that the file was correctly stored in the receiving file system. Intermediate mechanisms such as TCP checksums and sequence numbers just allow recovery in the common case, such as individual packet drop. These checks do ensure that the file was correctly transferred.

   (b) (5 Points) Describe how it has influenced the design of the original Internet protocols, with one specific example.

   Ans: The IP protocol contains very little mechanism, like no flow control, no checksum on data, no sequence numbering, to avoid "optimizations" being wired in that may not be useful for the various higher level protocols expected. Also, TCP does not mandate a keep-alive but expects higher-levels will provide if needed. Finally, one might cite UDP as providing transport-level demultiplexing but little else, expecting application protocols at the ends to provide the rest.

   (c) (5 points) Describe how the IP security standard (IPsec) fits into the overall end-to-end story.

   Ans: IPsec fits in the sense that an application can set up and use a specific IPsec security association, viewing the channel as for the application, just operating at IP level to do an end run on existing protocols, etc. One might also stretch to claim it fits in providing a higher-level of security over some exposed path, allowing the application to use a cheaper end-to-end encryption. However, it does not fit if used as part of the path in place of application-level security; the security than depends on trusting all the other components.

2. (15 points total) *Congestion*

   (a) (6 points) Al Gore, inventor of the Internet, lectured Monica Lewinsky that *congestion collapse* occurs from routers just getting tired from forwarding some many packets. However, even Monica didn't buy that line, or may be she slept through it. Describe what causes congestion collapse in reality, or is it just a fantasy of the liberal-biased media?

   Ans: Yes, congestion collapse occurs in reality, but from buffers filling with retransmissions. What collapses in the throughput, not the routers. Most packet drop is caused by buffer and queue limits. When hosts retransmit quickly in response to drop, they fill these buffers and queues with retransmits that cause more drops, including more drops of new packets. Combining that with reliable transport protocols that go back N in extreme conditions, the throughput can drop to almost nothing.

~~1~~
62

(b) (6 points) Describe the characteristics and/or limitations that make the Aloha protocol suitable for some radio networks and CSMA-CD suitable for the original Ethernet, and not vice versa, with suitable justification.

Ans: Aloha does not require senders to "hear" other senders so works with wider-area radio networks; each station just transmits when ready and hopes for the best. CSMA-CD requires the sender to listen before sending (carrier-sense) and during transmission (collision detection) to reduce the probability of collision. Thus, CSMA-CD provides much higher throughput of a network compared to Aloha but requires senders to be able to hear all other senders within its packet transmission time, thereby limiting the type, diameter/speed of network. Thus, Aloha is used when you cant use CSMA-CD but makes little sense to use when you can use CSMA-CD.

(c) (3 points) Bill Clinton decides to extend his managed healthcare initiative to cover the "health" of networks by legislating a new link-level flow control (LLFC) mechanism that is supposed to prevent congestion collapse. What are some concerns that the Republicans might legitimately raise about this approach.

Ans: Link-level flow control can cause a bottleneck node to push back and cause the buffers in upstream nodes to full, causing drop for packets other than those going to the real bottleneck. Moreover, without sufficient care, link-level FC can cause deadlock, where node A is waiting for buffers in B and B is waiting for buffers in A, and neither can free without sending and neither can send without free buffers. Finally, LLFC incurs network overhead, yet transport protocols such as TCP can sense the bottleneck rate and adapt, without this overhead. I.e. it's not needed in general and it costs. Finally, finally, what about multicast, Bill?

3. (15 Points total) *Bandwidth and Delay*

(a) (5 points) Suppose I need to move M megabytes of data from San Francisco to New York. One option is leasing a 1.5 Mbps Internet connection; Another is to write the data to 1.5 Megabyte floppy disks and hop on the plane for the 5 hour trip to New York. How big does M have to be for it to be faster to take the plane (stating whatever additional assumptions you need/make)?

Ans: 1.5Mbps times 5 hours (== 18,000 seconds) is 27 Gbits or 3,400 megabytes, so if M greater than 3400, taking the place might make sense, ignoring the time to write the floppies, and assuming zero overhead for the transfer over the leased line.

(b) (4 points) What factors could favor using the lease line over the plane trip, and vice versa, in reality.

Ans: Lease line might be better because you get some of the data sooner, you dont have to write the floppies, you dont risk the floppies being unreadable at the other end (or risk a long roundtrip if they are) and you might have need to update some of this data after the main transfer anyway. There may also be an issue of having to take the plane back, cost of plane relative to existing leased line, fitting with airplane schedules, etc. There may be also an issue of losing luggage on airplane, assuming you cant carry on 1800 floppies.

Favoring the plane, it would take a long time and extra cost to get the leased line in place if this was a one-off transfer. The plane would be cheaper unless the line was used a lot.

(c) (6 Points) Some routing experts advocate providing a bandwidth-delay product worth of buffering per port in the router, with delay referring to roundtrip time delay. For example, with 1.5 Mbps links and 25 millisecond round-trip time, the router should have roughly 4.7 kilobytes per port. Are these guys just trying to sell memory or what? Describe the legitimate argument for this amount of memory, if any, and how that relates to the memory required at the end points.

Ans: A host using a streaming transport protocol like TCP sending at the link rate can send a delay*bandwidth amount of data between the time a packet is dropped at a router on the path and it receives a corresponding drop indication (such as duplicate ack) from the receiver end. Thus, the router should be able to buffer that amount of data before expecting the sender to react to the packet drop. A similar argument holds for marking packets that go to receiver and then are returned. Conversely, with less buffering, the router may be forced to drop packets from the sender before it is possible for the sender to react to the congestion indication. The sender end should have at least the same amount of buffering so it can retransmit the dropped packet, and the receiver should too, so it can buffer the out-of-order packets until receiving the dropped packet, avoiding a complete go-back-N recovery.

4. (15 Points Total) *Routing*

(a) (6 points) Describe in brief how distance vectoring (DV) routing works.

Ans: Each node A computes its minimum distance d to each other node C based on the nodes it is directly connected to plus the similar information received from these other nodes. For instance, if neighbor B is 2 hops from C then A determines it is 3 hops from C, if no better route is advertized. It then exchanges this information with all its neighbors.

(b) (5 points) Describe the bad behavior that is inherent with DV routing with a specific example, and one solution that have been proposed/used to deal with it.

Ans: If a link fails, a node A may inadvertently select a new path to destination C that actually goes through itself because another node is advertizing a path to C that goes through A but is lower cost than anything else that A now has available. For instance, in the above, the link from A to B may fail, pmaking a path of D to C at cost of 4 look like the cheapest for A. However, D's path is based on going through A using its previous path of cost 3. On each exchange of routing information, the cost of this path goes up (because A's minimum cost path is now 5, then 7, etc.) However, it may take some time before the network realizes that C is unreachable or else has only a quite expensive route to it available. This problem is called inherent because DV routing does not distribute the topology information, just the distance vectors to destinations so nodes do not have the topology information to avoid this loopback problem.

One solution is for A to advertize an extremely high (infinite) cost to its route to C once this link fails and for some time after, forcing the path through it to look like a bad alternative to other nodes. I.e. their cost of the path going through A is then infinity plus their distance to A. This "holddown" accelerates finding a real path or discovering that these is no path to C.

(c) (4 points) Describe how link-state routing works, in essence, how it avoids the problems of DV routing, and at what costs.

Ans: Link-state routing exchanges the costs of links between all pairs of nodes N_i and N_j, providing each routing node with the costs and topology, allowing each to then compute the shortest paths to all destination based on its knowledge of the network toplogy and transit costs. In this approach, a link failure shows up as infinite cost on the failed link, causing other nodes to use alternative routes after the next update, if there is an alternative route. Here, each node knows the topology so does not make the mistake that DV does.

Costs are distributing the extra information about all links plus the cost of storing the topology and computing the shortest path to each destination from the given node, at each node in the network.

5. (1 Point) Name 3 top networking experts whose first name is David.

Ans: David Clark, David Mills, David Farber

The End — I'm clearly out of questions.