CS268 Exam Solutions

General comments:
1) If you would like a re-grade, submit in email a complete explanation of why your solution should be re-graded. Quote parts of your solution if necessary. In person requests for a re-grade will not be accepted.
2) The following solutions are more verbose than we would expect in an actual solution. However, some answers are too terse (e.g., "Link layer retransmissions are more efficient."). Excessively terse answers get partial credit.
3) In many cases, an advantage for one scheme is a disadvantage for another. If a solution contains the same general concept multiple times, we only credit it once.

1) End-to-End (20 pts)

(a) Retransmission of lost data can be done at the link, transport, and application layers. What are the pros and cons of doing it at each layer? Write no more than **three** phrases for each advantage/disadvantage.

Link layer:
+ Link layer retransmissions avoid having to retransmit over the entire path. Transport and application layer retransmissions must be sent over the entire path, thus incurring higher latency and wasting bandwidth.
+ Link layer retransmissions can hide non-congestion losses from higher layer protocols. For example, wireless link layer retransmissions can prevent error losses from being misinterpreted by TCP as congestion losses.
- End host to end host retransmissions are still necessary. A packet can be lost can after being successfully received at a router (due to congestion or errors). A route flap can send packets on a path that causes their TTL to expire before arriving at the destination. Some routers may not implement link layer retransmissions.
- Link layer retransmissions add some complexity and overhead. (2 pts)
- A packet may be retransmitted at multiple layers simultaneously. For example, a packet may be delayed before reaching a wireless link. The link layer protocol retransmits the packet, but the sender has already timed out and retransmits the packet. The link layer's retransmission is superfluous.

Transport layer:
+ Transport layer retransmissions avoid having to reimplement retransmission functionality in applications. Many applications can use the same transport protocol implemented in a library or in the operating system, reducing bugs, implementation time, etc.
+ Some solutions said that retransmissions and congestion control fit together naturally. This isn't entirely true. (2 pts)

Application layer:
+ Some applications can tolerate loss better than the increased delay of retransmissions. An example is an interactive audio or video application.

+ Data may be lost above the transport layer. An example of this is a buggy proxy cache. It could transfer data correctly from a server, corrupt the data, and then send the corrupt data to the client.

Grading: 4 points for each + or - above unless otherwise noted.

2) TCP and Route Fluttering (20 pts)

(a) Describe fast retransmission and fast recovery algorithms (pseudocode preferred).

(b) Consider a TCP flow implementing fast recovery. Assume the end-to-end delay of each packet is d1 in both directions. Suppose the 4^{th} packet transmitted by the sender experiences a delay of d2 > d1. Assume the first packet is transmitted at time t. What is the congestion window size at time t + 6*d1 + ε?

Note: Ignore the time it takes to send/receive a packet and assume that the receiver sends an ack as soon as it receives a packet, and the sender sends the next packet as soon as it gets an ack. Assume there are no packets losses.

a)

Fourth ACK for packet p received:

retransmit p; /* fast retransmission */

sstresh = cwnd/2; /* fast recovery */

cwnd = sstresh; /* actually it's sstresh + 3 to account for packets acked by DUPACKs */

b)

b.1) d1 = d2; cwnd = 8

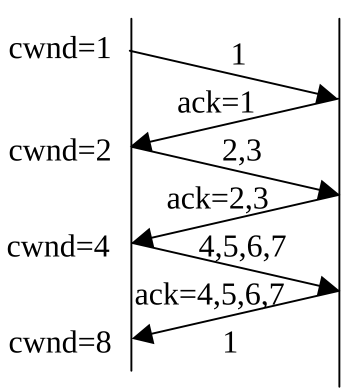

b.2) d1 < d2; cwnd = 2 or 5 (both answers were considered correct)

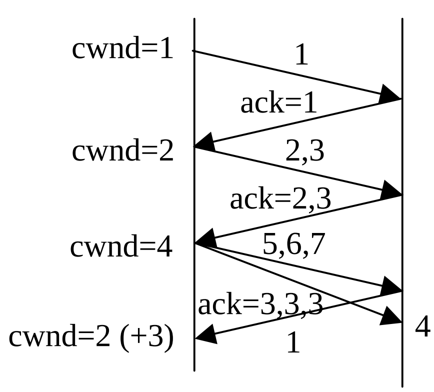

3) Differentiated and Integrated Services (20 pts)

How do premium service (Diffserv) and guaranteed service (Intserv) differ in their semantics and implementation?

Semantic differences:

- The guaranteed service is end-to-end; the premium service is ingress-to-egress (i.e., the premium service is defined on a domain basis)
- The guaranteed service provides **both** per-flow delay and bandwidth guarantees; the premium service provides only per-flow bandwidth guarantees

Implementation differences

- Premium service (and Diffserv in general) is deployed on a domain basis, and differentiate between edge and core routers
- With the guaranteed service, every router maintains per-flow state and performs per-flow packet processing and admission control; with the premium service, core routers maintain no per-flow state
- With the guaranteed service, the admission control is totally distributed; with the premium service, the admission control is centralized within each domain (i.e., it is done by the bandwidth broker)

4) Internet Addressing (20 pts)

(a) Classless Inter-Domain Routing (CIDR) allows the allocation of any power of 2 sized IP address range, instead of fixed size class A, B, and C subnets. Name two advantages of using CIDR.

CIDR allocates the address space 1) more efficiently and 2) reduces routing table size. For example, before CIDR, an organization with 2^10 hosts would get one class B (2^16) subnet or four class C (2^8 addresses) subnets. With one class B subnet, the organization would leave many addresses unused, thus reducing the overall efficiency of the IPv4 32 bit address space. With four class C subnets, those subnets could have different prefixes, resulting in four separate entries in routing tables. With CIDR, the organization could have one 10 bit address space, which it can use efficiently and which only occupies one entry in the routing table.

Grading: 4 points for each of the above.

(b) A mobile host travels to foreign networks while retaining its home IP address. To support this, mobile hosts advertise their home IP address to foreign routers, who would propagate this information to other routers during routing updates. Name two advantages and disadvantages of this scheme compared to Mobile IP.

New scheme:
+ This enables more efficient routing than Mobile IP. Mobile IP suffers the inefficiency of triangle routing.
+ Avoid the packet header overhead of encapsulation.
+ This is more fault tolerant than Mobile IP. The home agent in Mobile IP is a single point of failure.

+ The mobile host does not need to allocate an address in the foreign network.
- The new scheme would propagate routing updates slowly.
- The new scheme requires modifying routers. Routers would have to be modified to handle the new advertisements from hosts.
- If large numbers of mobile hosts use this scheme, routing tables would become very large. Every mobile host would create an entry in the routing table. Currently, this does not happen because all the hosts in a subnet can be aggregated into one routing table entry.
- The new scheme does not preserve location privacy. Anyone can use traceroute to determine which network the mobile host is connected to, which may give away the mobile host's geographic location.

Grading: 3 points for each of the above unless otherwise noted.

5) TCP Congestion Control (20 pts)

(a) A TCP flow's throughput is approximately $\frac{c}{r\sqrt{s}}$, where c is a constant, r is the round trip time, and s is the probability that a segment is lost. Assuming that the probability of packet loss is p, and each segment is split into n fragments, what is the TCP throughput?

$$\frac{c}{r\sqrt{1-(1-p)^n}}$$

For a TCP segment to be useful, all of its fragments must arrive at the receiver. The probability that one fragment arrives is $1-p$. The probability that n fragments arrive is $(1-p)^n$. Therefore, the probability that a segment is lost is $1-(1-p)^n$.

Grading: 8 points

(b) What are the pros and cons of explicit and implicit signals for congestion?

Implicit:
+ An Implicit congestion signal does not require router support.
+ Transport protocols always need to adapt to implicit congestion signals anyway. Explicit congestion signals can be lost, so a lack of an explicit signal does not imply a lack of congestion.
+ Some explicit signals require sending an extra packet to signal congestion. This can cause control traffic congestion collapse.
+ Some explicit signaling schemes set a congestion bit in packets. Selfish receivers can subvert this by not sending the bit back the sender.

Explicit:
+ An explicit congestion signal can be used for congestion avoidance. Routers can signal congestion when their queues grow long, but before they have to drop packets.

+ An explicit congestion signal requires lower latency to interpret than implicit congestion signals. For example, TCP requires 3 duplicate acknowledgements to determine that a packet has been lost.
+ An explicit congestion signal distinguishes congestion loss from other loss. Wireless networks may drop packets because of signal fading or interference. Route fluttering can send packets into oblivion. This sources of packet loss can be misinterpreted by an implicit congestion signal scheme.

Grading: 3 points for each of the above, unless otherwise noted.