CS 268 1st Term Exam
Date: March 22, 2006

Name:
SID:

Problem	Points
1	/ 20
2	/ 20
3	/ 10
4	/ 10
5	/ 20
6	/ 20
Total	

Please be **concise**.

1) E2E Arguments and Data Retransmission (20 pts)

Retransmission of lost data can be done at the link, transport, and application layers. What are the pros and cons of implementing retransmission at each layer?

Solution:

Link layer:

- Advantages: (1) improve performance, as receiver doesn't need to wait for sender to resend the packet; (2) more efficient, as packets are not traversing multiple links just to be dropped near the receiver, and then to be retransmitted again.

- Disadvantages: (1) doesn't eliminate the need for upper (e.g., transport) layers to perform retransmission; (2) may negatively interfere with retransmission of the upper-layer protocols (e.g., TCP).

Transport:

- Advantage: provide end-to-end retransmission; doesn't need to reply on the link layer to achieve correctness.
- Disadvantages: (1) inefficient and slow as a packet can only be retransmitted by the sender; (2) doesn't ensure application level reliability.

Application:

- Advantages: (1) satisfies the end-to-end arguments, as the application is in the best position to know what piece of data needs to be retransmitted, and what are the timeline constraints.
- Disadvantage: each application has to implement data retransmission.

2) Infinite Buffers (20 pts)

In a famous paper by Nagle, he argues that with earlier version of TCP (before Van Jacobson's work), even in a network with infinite buffer space at routers, the effective network throughput may go to zero.

(a) Explain why this could happen.

(b) Will Jacobson's work avoid this problem? Why or Why not?

(c) Will the DECBit algorithm avoid this problem? Why or Why not?

Solution:

(a) In the absence of congestion control, the queue will become larger and larger. This will lead to infinite delays, which in turn will trigger retransmission timeouts at the sender. As a result the sender will send more and more repair packets, up to the point that most of the packets in the network are repair packets. At this point the effective throughput goes to zero.

(b) Yes. When the link becomes congested, the delay will increase very fast, which will cause the retransmission timeout to expire even though the packets were not dropped (they are in the queue). As a result, the sender will reduce its sending rate (remember that in Van Jacobson's scheme the sender sets cwnd=1 after a retransmission timeout expires), which in turn cause the queue to drain.

(d) Yes. DECBit marks packets as soon as the link gets congested, and as a result the sender reduces the sending rate.

3) Router Architectures (10 pts)

Consider a router with N inputs and N outputs, where all inputs and outputs have the same capacity, and all speedups are **one**. Assume a uniform random load, i.e., a packet arriving at a particular input is forwarded to an output with equal probability. Under this load, which architecture do you expect to achieve a higher throughput: output-queued routers, or input-queued routers? Explain your answer.

Solution: Input-queued routers. If all speedups were one, it shouldn't matter whether the outputs have buffers or not. In other words, in this case output-queued routers have no better performance than a router with no buffers! On the other hand, input-queued routers can buffer packets if there is contention at the output interfaces, and then send them later.

4) Multicast (10 pts)

Consider a binary balanced tree with L levels, where the root (level 0) wishes to send a packet to all leaves (at level L). Let M denote the total number of packets *forwarded* in the tree for every packet sent by the sender. A packet traversing one edge of the tree is counted exactly once.

(a) Compute M assuming that all routers implement multicast.

(b) Compute M assuming that none of the routers implement multicast.

(c) Compute M assuming that only routers at level k < L are multicast-aware. What is the decrease in M, as compared to (b), when k = L/2, and L goes to infinity?

Note: A multicast aware router *A* will send a unicast packet to every leaf in the tree routed at *A*.

Solution

(a) $2^{L+1} - 2$

(b) $L*2^L$

(c) $L*2^L + k*(2^k – 2^L)$ → $(L*2^L + k*(2^k – 2^L))/ (L*2^L)$ when k = L/2 and L goes to infinity approaches 1/2

5) 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.

(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 two disadvantages of this scheme compared to Mobile IP.

Solution

(a) Advantages:

- Allows finer grain address allocation → improves address allocation efficiency.
- Improves address aggregation → reduces the router table size.

(c) Advantages:

- Avoids triangular routing problem.
- Doesn't require home agent; more resilient because no single point of failures.

Disadvantages:

- Breaks address aggregation.
- Need to modify routing protocol.
- Compromise anonymity.

6) TCP Congestion Control (20 pts)

Consider a TCP version that uses two disjoint paths instead of one between source and destination. Let p_1, p_2 and RTT_1, RTT_2 be the packet loss rates and round-trip times along the two paths.

(a) Describe two challenges you need to address in adapting the TCP to use two disjoint paths, instead of one.

(b) Compute the TCP throughput when using both paths. Give the condition (e.g., a relation between p_1, p_2 RTT_1, and RTT_2) under which the TCP throughput is higher when using two paths instead of one.

Note: The throughput of a TCP flow along a path with loss rate p and round-trip time RTT is $c/(RTT*p^{1/2})$, where c is a constant.

Solution:

Assume the source estimates a single RTT. (Solutions assuming that the sender estimates an RTT for each path also received full credit.)

(a) Correctly estimate the end-to-end RTT, and handle packet reordering at the receiver due to differences in the RTT along the two paths.

(b) When using both paths, the TCP throughput is $(c/p_1^{1/2} + c/p_2^{1/2})/\max(RTT_1, RTT_2)$.
When using one path the maximum TCP throughput is
$\max(c/(RTT_1*p_1^{1/2}), c/(RTT_2*p_2^{1/2}))$