Exam #2 for Computer Networks (CIS 6930) Fall 2000

>>> *SOLUTION* <<<

Welcome to Exam #2 in Computer Networks (CIS 6930). You have 75 minutes. Read each problem carefully. There are six required problems (each worth 16 points - yes, you get 4 points overall for "free") and one extra credit problem worth 5 points. You may have with you a calculator, pencils, eraser, blank paper, lucky rabbit's foot, and one 8.5 x 11 inch "formula sheet". On this formula sheet you may have anything you want (definitions, formulas, etc.) *handwritten by you*. You may use both sides. Computer generated text, photocopies, and scans are not allowed on this sheet. Please submit your formula sheet with your exam. Please start each numbered problem on a new sheet of paper and do not write on the back of the sheets (I do not care about saving paper!). Submit everything in problem order. No sharing of calculators. Good luck and be sure to show your work. <u>Be PRECISE and CONCISE in all of your answers</u>.

<u>Problem #1</u> (20 minutes)

Describe IEEE 802.3 Ethernet. You should describe the intended application of Ethernet, CSMA/CD, frame format, addressing, frame lengths (what are they and why), topology, history, and scalability to higher speeds and larger spans.

```
IEEE 802.3 Ethernet is LAN technology originally intended to support data
networking in a single building with 1024 attached stations, about 2
kilometer span, and 10-Mbps data rate.  Ethernet uses Carrier Sense
Multiple Access / Collision Detection as its MAC protocol to share a common
medium.  The CSMA/CD protocol is:

(1) if (medium is idle) then xmit
(2) if (medium is busy) then wait until idle and when idle xmit immediately
(3) if (detect a collision when xmitting) then stop xmitting immediately
and back-off a random period of time
(4) goto (1)

The Binary Exponential Backoff (BEB) protocol for determining back-off time
is:

while (attempts < 10)
  k = min(attempts, 10)
  k = rand(0, 2**k)
  delay = r * slot_time

The frame format is as follows:

+----------+----+----+-----+------//-+-----+
| preamble | DA | SA | Len | data     | FCS|
+----------+----+----+-----+------//-+-----+
    8 bytes  6    6     2     var       4

Where the minimum packet length is 64 bytes based on a 2 * t_pr minimum
length for CD before transmission complete in a 2 kilometer span network.
The maximum packet length is 1500 bytes based on a historic 1K block size
and that putting 1500 bytes of memory on a card was the maximum feasible.
Addresses (SA) are uniquely defined in each adapter and destination
addresses can be either to a specific address or to broadcast (all stations
- address of all 1's).
```

The topology of Ethernet is logical bus and has evolved from a physical direct-wired bus (10BASE5 and 10BASE2) to star-wired (e.g., 10BASE-T). Ethernet was invented by Metcalfe in 1976 motivated by Abramson's ALOHA packet radio. From 1976, Ethernet has evolved from coax 10BASE5 and 10BASE2 to UTP 10BASE-T and 100BASE-T. Scalability of CSMA/CD is limited by the need for a minimum frame length to be 2*t_pr, thus future 10-Gbps Ethernet does not specify CSMA/CD.

Problem #2 (10 minutes)

Describe layer-2 bridging. Describe the bridging table and forwarding/learning algorithm. Explain why a spanning tree is needed. Give the three names used for a bridge and describe how each name is appropriate.

Layer-2 bridging serves as a MAC level filter between LANs with like MACs. The bridging table contains MAC addresses of all observed stations in the network and a direction for each address (local = "behind this port" or not local = "ahead of this port"). The key algorithm is:

```
(1) Receive a frame (interface is promiscuous)
(2) If (DA in table) and (DA local) then do not forward else forward
(3) If (SA in table) check/update direction
(4) If (SA not in table) add to table with direction
```

A spanning tree is needed to prevent loops. Loops would confuse the learning algorithm (i.e., the bridge would never learn on "what side" is a station) and cause infinite propagation of frames in the loop. Three names are: "Transparent bridge" because the bridge is "transparent" to end stations (i.e., no changes are needed in the end stations). "Learning bridge" because the bridge learns its forwarding table as described above. And, "spanning tree bridge" because the network of bridges must form a spanning tree for correct operation.

Problem #3 (10 minutes)

Assume a 64-kbps, 2500 kilometer full-duplex link with a single sender and receiver. The sender always has packets queued to send. Packets are always 250 bytes in length. Medium propagation is 5 nanoseconds/meter. Bit errors are independent. Solve for link utilization, *U*, for the following cases:

We will assume that t_ack and t_proc are negligible since no parameters are given to determine these values.

a) Stop-and-wait flow control, Pr[bit error] = 0

```
U = t_fr/(t_fr + 2*t_pr).

t_fr = (250 * 8) / 64000 = 31.25 millisec
t_pr = 2500000 * 5 * 10**-9 = 12.5 millisec

U = 31.25 / (31.25 + 2 * 12.5) = 55.56 %
```

b) Stop-and-wait flow and error control, Pr[bit error] = 0.0001

```
p = Pr[frame loss] = 1 - (1 - Pr[bit error])**(250*8) = 0.1813

U = (1 - p) * t_fr / (t_fr + 2*t_pr) = 45.48%
```

c) Sliding window flow control with window size (*N*) of 5, Pr[bit error] = 0

```
Since N*t_fr > t_fr + 2*t_pr and no errors, U = 100%
```

d) Sliding window flow control, Go-Back-N (GBN) error control, window size of 5, and Pr[bit error] = 0.0001

```
Since N*t_fr > t_fr + 2*t_pr and errors,
U = (t_fr * (1 - p)) / (t_fr + (2 * t_pr * p)) = 71.50%
```

Problem #4 (10 minutes)

What is the purpose of ARP? Describe how ARP works.

The purpose of Address Resolution Protocol is to resolve IP and MAC addresses for IP hosts on a LAN. Each host in an IP network contains an ARP cache - a one-to-one table of IP addresses and associated MAC addresses. When the host queues a packet for transmit, the ARP cache is used to "fill in" the MAC destination address. If the ARP cache does not contain an entry for the destination IP address, ARP broadcasts an ARP request to all hosts on the LAN. Only the host with the targeted IP address, or the proxy for the address (typically a router), will respond. This response contains the source MAC address of the host handling the target IP address and now the ARP cache can be updated for this IP address. Entries in the ARP cache are aged-out if not used.

Problem #5 (15 minutes)

Describe how an IP router works.

An IP router uses a routing table consisting of net_ids and port numbers to forward IP datagrams, or packets, to the destination net_id through the "best" port. A port is a physical connection to a network (typically, a network is a LAN). The best port is the one that leads to the shortest path. The routing table must contain all the net_ids in the autonomous system. A packet intended for a net_id not in the routing table is sent out on a default gateway (e.g., it sent to the Internet). The routing table is generated via a distributed routing protocol such as RIP or OSPF that constantly computes least-cost paths between networks. A least-cost path can be lowest delay (i.e., adaptive routing) or simply lowest hop count. An incoming packet to a router has its MAC addresses swapped to the destination address of the next router, or the host if routing to a target network. The ARP cache at each port resolves host_id to MAC address for target hosts or resolves net_id for routers proxying for other networks. Thus, routing tables consist only of net_ids and port numbers reducing the overall size (i.e., compared to if the routing tables contained host_id information in addition to net_id information).

Problem #6 (10 minutes)

Describe how TCP adaptively determines its packet-is-lost retry time-out value (RTO). What are the consequences of a too long or too short time-out value? Which is “better”, too long or too short of a time-out?

Van Jacobson’s method computes a smoothed RTT estimate (SRTT), a smoothed error estimate (SERR), and a smoothed deviation estimate (SDEV). From this RTO is directly computed as SRTT + (f * SDEV). The time-out is computed this way to determine a “good” time-out for both short/low-delay-variance and long/large-delay-variance connections. It would be impossible to fix a single time-out value for all possible TCP/IP connections. A time-out that is too short causes duplicate packet sends, one that is too long causes excessive waiting when a recovery should be taking place. In both of these cases, efficiency is affected. A too long value is “better” since it hurts only the sending host whereas a too short value puts excess, redudant load on the network (i.e., is bad for the community).

Extra Credit Problem

Give one good suggestion on how to improve this course. This is not an easy question because your suggestion must not entail something that requires significant amounts of additional work, or reduces the content of the course. Thus, suggestions such as “hold help sessions every Saturday from 2pm to 5pm” and “eliminate the second test” are not good suggestions (and will gain you no points).

Various answers are possible.