

**NAME:**

Login name:

---

**Computer Science 461**  
**Midterm Exam**  
**March 10, 2010**  
**3:00-4:20pm**

---

This test has five (5) questions. Put your name on *every page*, and write out and sign the Honor Code pledge before turning in the test.

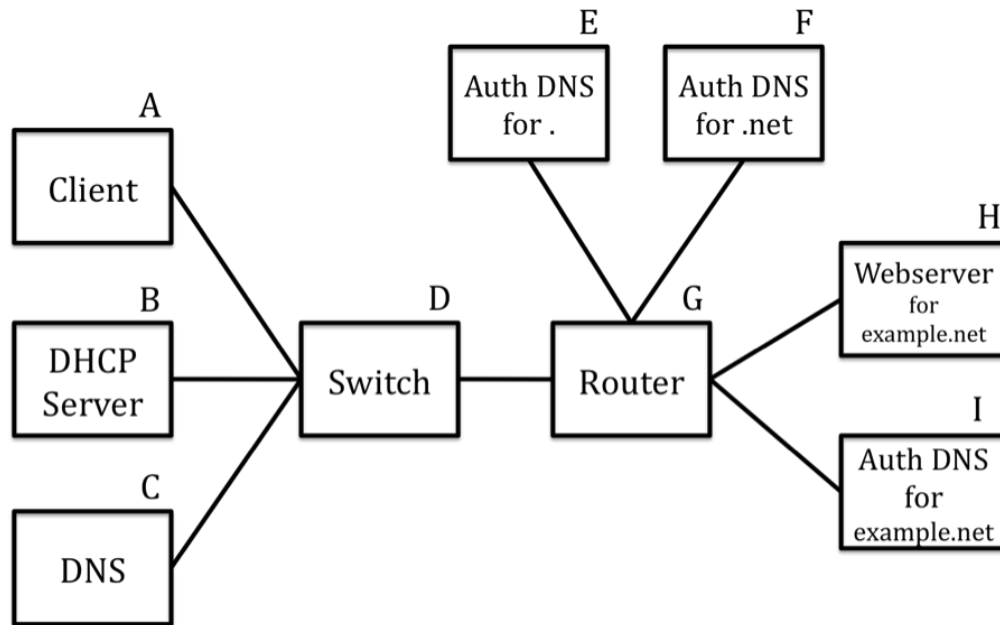
Please look through all of the questions at the beginning to help in pacing yourself for the exam. The exam has 100 points and lasts for 80 minutes, so the number of minutes spent per question should be *less* than its point value. You should spend no more than 15 minutes per question. Show your work for all problems. Partial credit will often be given.

"I pledge my honor that I have not violated the Honor Code during this examination."

| <b>Question</b> | <b>Score</b> |
|-----------------|--------------|
| <b>1</b>        | <b>/ 19</b>  |
| <b>2</b>        | <b>/ 19</b>  |
| <b>3</b>        | <b>/ 16</b>  |
| <b>4</b>        | <b>/ 29</b>  |
| <b>5</b>        | <b>/ 17</b>  |
| <b>Total</b>    | <b>/ 100</b> |

**QUESTION 1 : Protocols and Addressing (19 POINTS)**

In the topology shown below, a client A (on the left side) joins a network and wants to download the webpage for `http://example.net/` from the webserver H (on the right side). All wires/links are Ethernet segments. In this question, you will be asked what packets are on the wire (and the addresses in those packets' headers) for this to occur. When asked to provide the addresses of machines, you can refer to them abstractly: IP-H is the IP address of the webserver H, MAC-H is the MAC address of H. If a packet is being sent to a broadcast address, use "BR" as the destination (e.g., IP-BR or MAC-BR).



Assume the following:

- All the machines, except the client, already have their ARP caches full.
- DNS resolver C has an empty DNS cache.
- `http://example.net/index.html` fits in a single packet.
- Client just joined the network and so all of its caches are empty (DNS, ARP, DHCP)

When you are asked to write a packet with its headers, you should include both network and link-layer addresses. So a DNS query from client A to server C and the response should look like the following. Note that other packets may be involved with a full DNS lookup that we are not showing in this example.

| Order | Link Src | Link Dst | Network Src | Network Dst | Protocol Description |
|-------|----------|----------|-------------|-------------|----------------------|
| 1     | MAC-A    | MAC-C    | IP-A        | IP-C        | DNS query            |
| 2     | MAC-C    | MAC-A    | IP-C        | IP-A        | DNS resp.            |

**In the following questions, you may not need to use all the lines we provide.**

1. [4 points] As it just joined the network, the first thing the client needs to do is get an IP address and other related network state. In the next lines, write the packets on the network involved in that step, including their header information.

| Order | Link Src | Link Dst | Network Src | Network Dst | Protocol Description |
|-------|----------|----------|-------------|-------------|----------------------|
| 1     |          |          |             |             |                      |
| 2     |          |          |             |             |                      |
| 3     |          |          |             |             |                      |
| 4     |          |          |             |             |                      |
| 5     |          |          |             |             |                      |
| 6     |          |          |             |             |                      |

2. [10 points] Now that the client has an IP address for itself, it wants to resolve the IP address of example.com, which will be IP\_H. Write down all the packets involved in this phase. Remember that the client just joined the network.

If a packet's headers change during transmission (i.e., the IP or MAC address get changed), you should write the packet on two separate lines, corresponding to the two set of headers it sees during transmission. For the description, write the most specific protocol, i.e., "DNS query", not "UDP".

| Order | Link Src | Link Dst | Network Src | Network Dst | Protocol Description |
|-------|----------|----------|-------------|-------------|----------------------|
| 1     |          |          |             |             |                      |
| 2     |          |          |             |             |                      |
| 3     |          |          |             |             |                      |
| 4     |          |          |             |             |                      |
| 5     |          |          |             |             |                      |
| 6     |          |          |             |             |                      |
| 7     |          |          |             |             |                      |
| 8     |          |          |             |             |                      |
| 9     |          |          |             |             |                      |
| 10    |          |          |             |             |                      |
| 11    |          |          |             |             |                      |
| 12    |          |          |             |             |                      |
| 13    |          |          |             |             |                      |
| 14    |          |          |             |             |                      |
| 15    |          |          |             |             |                      |
| 16    |          |          |             |             |                      |
| 17    |          |          |             |             |                      |


3. [5 points] Now that the client has IP\_H, it wants to download the web page <http://example.net/index.html>. (Recall that the webpage fits in one data packet.) Write down all *network-layer* packets that are sent on the wire in order to complete this process. Remember to clean up any transfers.

For the “protocol description”, if the packet is associated with transport layer, you should write something like “TCP SYN”. If it’s inherently an application-level message (i.e., data sent over a TCP connection), you should write something like “HTTP request”. In this example, you should assume that you are running a “basic” TCP implementation, so it’s not doing any fancier optimizations like delayed ACKs or piggybacking.

| Order | Network Src | Network Dst | Description |
|-------|-------------|-------------|-------------|
| 1     |             |             |             |
| 2     |             |             |             |
| 3     |             |             |             |
| 4     |             |             |             |
| 5     |             |             |             |
| 6     |             |             |             |
| 7     |             |             |             |
| 8     |             |             |             |
| 9     |             |             |             |
| 10    |             |             |             |
| 11    |             |             |             |
| 12    |             |             |             |

## ***QUESTION 2 : TCP and Congestion Control (19 POINTS)***

Consider the following graph of TCP throughput (**NOT DRAWN TO SCALE**), where the y-axis describes the TCP window size of the sender. We will later ask you to describe what happens on the right side of the graph as the sender continues to transmit.



1. [ 3 points ] The window size of the TCP sender decreases at the points marked by A, C, and D on the graph, which happens when a packet belonging to the stream is lost. Name the event which occurs that causes the sender to decrease its window.

A:

C:

D:

2. [ 4 points ] Assume that the network has an MSS of 1000 bytes. If point A occurs 1 second after the sender begins its transfer, and the sender has written 15,000 bytes to the network by that time, what is the round-trip-time (RTT) of the network? Assume at time 0 the sender attempts to open the connection. Also assume that the sender can “write” a full window’s worth of data instantaneously, so the only latency you need to worry about is the actual propagation delay of the network.

3. [ 4 points ] What is the sender’s window size (in bytes) at point B?

4. [ 4 points ] If point C occurs 2 seconds after point B, what is the sender’s window size (in bytes) at point C?

5. [ 4 points ] Assume that the window size at point D is 16,000 bytes, at which time it drops to 1,000 bytes. How much time will it take for the sender to return to a window size of 16,000 bytes (assuming that there is no further packet loss in the network)? You can express your answer in terms of the number of RTTs, as opposed to the number of seconds.

### ***QUESTION 3: Socket Programming (16 points)***

You are implementing a simple application to transfer a file from the client to the server using TCP.

1. [ 5 points ] In the following, write the basic system calls you need to use in order to transfer the file. Assume that the client already has the IP address of the server. You don't need to show arguments, i.e., name resolution could have just been written as `gethostbyname` (not `struct hostent * gethostbyname(const char *name)`) and you don't need to show any other programmatic logic. Assume that the file can be written/read in a single call to `send/recv`.

In the below table, provide a set of system calls in an appropriate order. Each line should have a single call at the client or server, not both. Some calls at the client and server can happen concurrently; just list them in a reasonable order (i.e., the write at the client should occur before the read at the server). Note: You might not need to use all the rows supplied. Remember to cleanup after you transmit the file.

|           | <i>Client</i> | <i>Server</i> |
|-----------|---------------|---------------|
| <b>1</b>  |               |               |
| <b>2</b>  |               |               |
| <b>3</b>  |               |               |
| <b>4</b>  |               |               |
| <b>5</b>  |               |               |
| <b>6</b>  |               |               |
| <b>7</b>  |               |               |
| <b>8</b>  |               |               |
| <b>9</b>  |               |               |
| <b>10</b> |               |               |
| <b>11</b> |               |               |
| <b>12</b> |               |               |
| <b>13</b> |               |               |
| <b>14</b> |               |               |
| <b>15</b> |               |               |

2. [11 points] Now consider the pseudocode of the client's `send` call in a bit more depth. Write some pseudocode (it doesn't need to compile, but should look C-like), how you would transmit a 100,000B file. In this code snippet, you *should* use arguments and return values to express the logic flow.

Your code should only be a few lines long; we're only interested in the logic immediately surrounding your use of the `send` system call.

```
// buf points to a buffer of size len
//   that you want to transmit in its entirety
// fd is an open file descriptor connected to server
// RETURN 0 ON SUCCESS, -1 ON FAILURE
```

```
int sendall (int fd, char *buf, int len) {
```

```
}
```



## ***QUESTION 4: Short answers (29 points)***

### 1. [6 points] IP Address Aggregation

Suppose you are network operator for an ISP, and your ISP has two customer networks (MIT and Princeton). For the sake of simplicity, let's assume that MIT only has 3 hosts with the following IP addresses, and Princeton only has 2 hosts:

MIT: 18.0.0.1, 18.0.0.254, 18.255.255.254

Princeton: 128.112.0.1, and 128.112.255.254

You've heard that memory in router line cards is expensive, so you'd like to use as few entries as possible in your line cards. Your colleague at work says he has a solution that only requires two forwarding-table entries.

(a) Give the IP prefix notation that your colleague uses to represent these two networks.

(b) After you install the forwarding table entry your colleague suggests, you notice that you still see traffic for host 18.255.255.255, but you've stopped seeing any traffic for hosts 18.0.0.1 and 18.0.0.255. So, you come up with a new set of forwarding-table entries – three this time – that, once installed, causes you to receive traffic for all five hosts. What are these three IP prefixes?

(c) Give one possible explanation why this change in part (b) is needed.

2. [ 4 points ] Ethernet and IP

(a) Suppose Ethernet was the only existing LAN technology, so every host in the Internet was part of a local Ethernet and thus had a globally-unique Ethernet address. Would you recommend getting rid of IP addresses by simply using Ethernet addresses instead of IP addresses? Why or why not?

(b) What about the other way around, why do we not simply assign IP addresses to network adaptors, instead of dealing with both Ethernet and IP addresses?

3. [ 8 points ] Using TCP, a sender has sent out packets 1 through 20. The sender receives an ACK for packet 10, then it receives three ACKs for packet 11 (i.e., if the final byte of packet 11 was byte no. 5000, then these ACKs were for byte 5001). Given this result:

(a) Which of the 20 packets can the sender assume are lost?


(b) Which of the 20 packets can the sender assume were definitely received?

Now, the same sender resends the next missing packet and receives an ACK for packet 16 in response. Based on this information:

(c) Which packets of the original 20 can the sender assume are still lost, if any?

(d) Which of the original 20 can the sender assume were definitely received?

4. [ 5 points ] Assume you have a network with three hosts – A, B, and C – all directly connected to the same switch. Host A first attempts to open a TCP connection to B at time 0, during which time he adds an entry to his ARP cache with a TTL of 60 seconds. Host B begins transferring a large file to host A. A new host D connects to the network, but is configured with the same IP address as already being used by A.



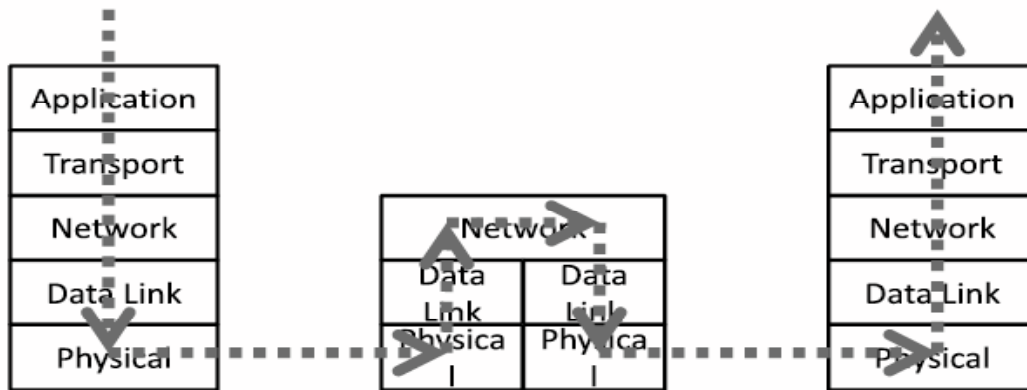
At time 55 seconds, an application running on host D opens a TCP connection to an application on C. A few seconds later, the TCP transfer between hosts A and B breaks. What were the steps that led to this connection breaking?

5. [ 6 points ] What transport layer protocol (TCP or UDP) you would use if your application needed one of the following properties and then give one succinct reason why. Be specific.

- (a) Speed
- (b) Long messages
- (c) Short-duration interactions
- (d) Fairness between competing flows
- (e) Flooding
- (f) Flexibility

**QUESTION 5: Multiple choice (17 points)**

1. What is the name of the forwarding device in the middle of the diagram below?  
(Circle ONE answer – 1 point)



- (a) Switch  
(b) Router  
(c) Hub  
(d) Repeater
2. When a UDP segment arrives at a host, in order to direct the segment to the appropriate socket, the operating system's network stack uses the following fields:  
(circle ALL that are true – 2 points)
- (a) the source IP address.  
(b) the destination IP address  
(c) the source port number  
(d) the destination port number
3. When a TCP segment belonging to an existing connection arrives at a host, in order to direct the segment to the appropriate socket the operating system's network stack uses the following fields (circle ALL that are true – 2 points)
- (a) the source IP address.  
(b) the destination IP address  
(c) the source port number  
(d) the destination port number

4. Which of the following link protocols best describes Ethernet? (circle ONE – 1 point)
- (a) frequency-division multiple access (FDMA)
  - (b) time-division multiple access (TDMA)
  - (c) token-passing
  - (d) carrier sense multiple access (CSMA)
5. If you are building a network in which only one host will be sending traffic, which of the following link protocols would give the best throughput? (circle ONE – 1 point)
- (a) frequency-division multiple access (FDMA)
  - (b) time-division multiple access (TDMA)
  - (c) token-passing
  - (d) carrier sense multiple access (CSMA)
6. Which of the following are advantages that RED has over drop-tail queuing? (circle ALL that are true – 3 points)
- (a) has fewer burst losses
  - (b) has shorter queuing delays
  - (c) ensures roughly equal throughput for all flows
  - (d) ensures roughly equal throughput for all TCP flows
7. Which of the following are true about Ethernet networks? (circle ALL true – 3 points)
- (a) Ethernet frames have a minimum size to ensure good utilization of the network (i.e., senders could otherwise have a large overhead to send a small piece of data).
  - (b) Bridges provide greater scalability for Ethernet networks than hubs.
  - (c) Hosts on different segments of a switched Ethernet network will only see packets to hosts on their segments or to the broadcast address.
  - (d) When an Ethernet sender detects that the media is idle, it sends a jam signal onto the media to tell other devices not to transmit, and then it sends its packet.
8. Which is true of the UNIX Socket API? (circle ALL that are true – 3 points)
- (a) If a chunk of data written to a TCP socket is smaller than the MSS, it will all be sent in the same packet.
  - (b) Writing a single message to a UDP socket can result in the network stack sending multiple packets.
  - (c) The socket allocated by the accept() system call is assigned a new port number
  - (d) You can call connect() on a UDP socket
9. [1 point] Describe one thing you would like to see changed in the class. (Only wrong answer is "nothing")