15-441:Networking

Homework 3

Homework 3 is due Monday, April 10th at 16:59:59. Submission will be electronic, as for the first 2 homeworks.

Problem 1: Packet Headers

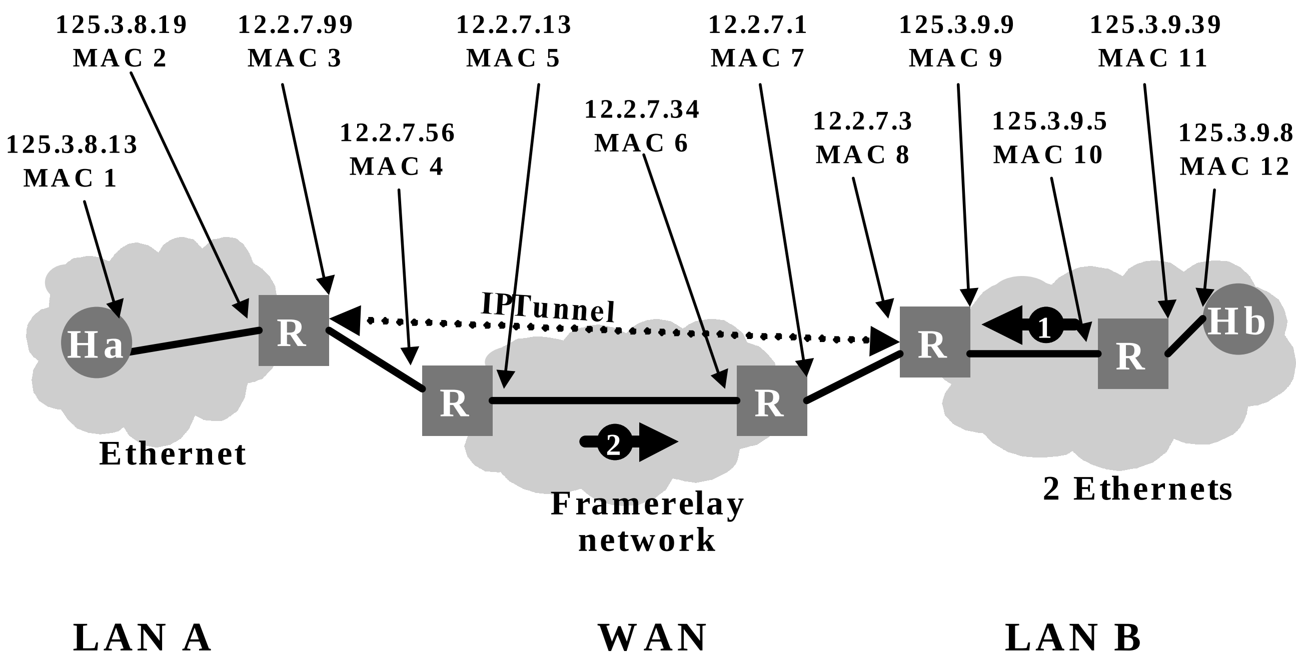

A corporation has two local area networks connected over a wide area network using an IP tunnel, as is shown in the picture above. The two edge routers that connect the LANs to the WAN forward any packets destined for the other LAN through the tunnel. Tunneling is used so that the corporation can more easily use non-standard protocol features in its corporate network. The first LAN is a simple Ethernet while the second LAN consists of two Ethernets connected by an internal router. Part of the WAN runs through a Frame Relay network. In the figure, R stands for Router and H for Host. The Frame Relay switches are not shown in the picture, and Ethernet hubs and switches are also not shown.

We want to look in more detail at the headers of packets being exchanged by a web browser executing on the client host H_a in LAN A and the web server H_b executing in LAN B.

a) Please show the type and the order of the headers of the packets being exchanged by the web browser and server if they are captured at locations "1" and "2" as indicated on the network picture. When showing the "type" of header, be specific (e.g., Don't just write 'transport', but rather indicate 'TCP' or 'UDP'. Don't just write 'datalink', but rather indicate 'Frame Relay', or 'Ethernet' or something else.) Include all headers of the following protocol layers: network layer, application layer, datalink layer, and transport layer. Use the packet outlines shown below, i.e., fill in the header types in the right box. Note that there may be more boxes than you need. The top of each packet outline corresponds to the first byte in the packet while the bottom corresponds to the last byte.

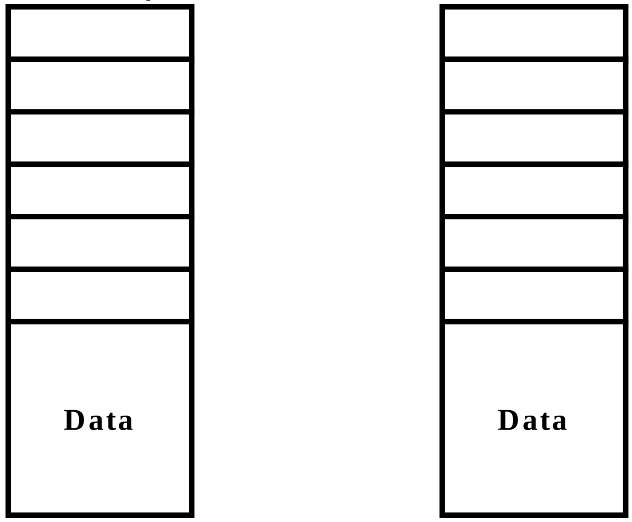

Packet 1: Hb -> Ha **Packet 2: Ha -> Hb**

b) For the packet traveling in LAN B from LAN B to LAN A, identify the primary fields of the protocol headers, meaning transport-layer source and destination port numbers, network-layer source and destination IP addresses, and link-layer addresses or connection identifiers. There is no need to fill in anything for application headers. Use the IP and datalink address information shown in the figure. The URL that the browser is trying to retrieve is ”http://www.internal.supercom/index.hmtl”. If you do not have enough information to specify a field, fill in an ”X” and briefly explain. If a field does not exist, fill in ”-” and explain briefly.

Hdr Type	**Src**	**Dest**
	Data	

c) The same question for the packet traveling from LAN A to LAN B through the Frame Relay network.

Hdr Type	Src	Dest
	Data	

Problem 2: DNS

The Andrew Linux machines provide a program `dig` that allows you to query Domain Name Service (DNS) servers around the Internet (some documentation is available if you type `man dig`). When running dig for the purposes of this question, you should use the following format:

dig +norecurse @name.of.dns.server record-type domain-name where

- *name.of.dns.server* is the hostname of the DNS server you wish to query, such as A.ROOT-SERVERS.NET.
- *record-type* is the type of DNS record you wish to retrieve, such as ANY, MX, HINFO, A, or SOA.
- *domain-name* is the name of the host or domain you seek information on.

The DNS is a distributed architecture that uses hierarchical delegation. At the top of the system are the "root" name servers, who know which DNS server is responsible for each second-level domain (such as CMU.EDU). If you send a root server a query for a particular machine, you will receive a reply listing the servers that have been delegated authority for that machine's second-level domain. It is common for a large domain such as CMU.EDU to further delegate to "departmental" or workgroup DNS servers, which you can discover by querying the second-level servers.

(a) In order to discover the chain of delegation in use at CMU, run a series of NS queries for UX15.SP.CS.CMU.EDU. You may start with any of the root servers, and you should continue your sequence of queries until you stop getting new delegations (in some domains, this is indicated by a DNS server returning you a delegation pointing to itself, and in other domains this is indicated by a DNS server returning you a SOA record instead).

Delegation chain for: AOL.COM

```
Server queried                 NS delgations to
--------------                 ----------------
A.ROOT-SERVERS.NET             A.GTLD-SERVERS.NET, K.GTLD-SERVERS.NET
K.GTLD-SERVERS.NET             DNS-01.NS.AOL.COM, DNS-02.NS.AOL.COM
DNS-01.NS.AOL.COM              DNS-01.NS.AOL.COM, DNS-02.NS.AOL.COM
```

This was produced by running the following commands:

% dig +norecurse @a.root-servers.net NS aol.com

% dig +norecurse @k.gtld-servers.net NS aol.com

% dig +norecurse @dns-01.ns.aol.com NS aol.com

Generate the delegation chain for UX15.SP.CS.CMU.EDU. Present your results in the table form shown above. Each NS query will typically return two or more answers; choose among them at random. If you query a server and get a timeout, choose an alternate server.

(b) The DNS is also used to translate IP addresses into hostnames. Again, the database is distributed in a hierarchical fashion, with a wrinkle. The most-specific part of a domain name is on the left (i.e., UX15 in UX15.SP.CS.CMU.EDU), but the reverse is true of IP addresses (i.e., in 128.2.203.134, 128 is "top-level", 128.2 is CMU.EDU, and 128.2.203 belongs to CS.CMU.EDU). Thus, address-to-name mapping is handled by reversing the bytes of the IP address and making queries in a special domain. To turn 128.2.203.134 into a hostname, various servers are sent queries seeking PTR records for 134.203.2.128.in-addr.arpa. The first query would be:

% dig @a.root-servers.net PTR 134.203.2.128.in-addr.arpa

You will know you're done when your query gives you back a PTR record in addition to (or instead of) NS records.

Fill in a table like the one above showing a query chain for the IP address 64.91.109.37.

```
Server queried                 NS delegations (or PTR record)
--------------                 ----------------
...                            ...
```

Problem 3: Mars Exploration

You are part of an exploratory mission to Mars. When, after a long trip, the mission arrives on Mars it discovers the remains of an ancient civilization. Shockingly, it appears that the previous inhabitants of Mars used 4Mb/s token ring for their computer infrastructure. Sadly, however, while your team easily locates computer interface cards, cabling, and hubs, it appears that all the transmit tokens are missing.

When you try to report your findings (including beautiful color pictures of the writing on the Martian token-ring hubs) to Mission Control on Earth, you discover that transmission is extremely slow. It takes almost three hours (10,000 seconds, to be exact) to send just one picture (of 1000x1000 pixels, 8 bits per pixel). The radio link transmits at 800 KBit/second, but the average bit rate you are achieving to Mission Control is only a fraction of this speed.

Your mission colleagues ask you to apply your 15-441 expertise to diagnose the problem. After some digging, you discover that the communication link uses a very simple point-to-point data-link protocol: it uses flow control and assumes the transmission channel is reliable. The packet size is 1000 bytes. You also discover that the radio transmission system was originally built for the Apollo 11 mission to the moon. You immediately suspect that the engineers who designed your Mars spacecraft did not take 15-441!

	Moon	Mars
Distance to Earth (km)	390,000	390,000,000
Diameter (km)	3475	6000
Distance to Sun (km)	150 M	230 M

A: Do your best to explain the problem. Please be precise (i.e., give numbers). We have provided some information (which may or may not be relevant) above about distances at the time of the mission. Also, the speed of radio waves in a vacuum is 300,000 km/second.

B: How would you fix the problem? Again, be precise. How practical do you think this is?

C: Once you have fixed the throughput problem, mission control is reporting corrupted pictures. It seems that a radio transmission system designed to work from the moon does not quite cut it from Mars. What changes would you make to the datalink protocol to recover from the transmission errors?

Problem 4: TCP Handshake

After leaving CMU you are hired by a new client/server middleware startup called Super-Tricky.com. After you've been there a few days, your manager drops by your office to ask if you to evaluate a cool new idea the company has been considering.

The problem is that applications based on their middleware frequently need results from multiple servers. For an application to contact one server and get a response TCP requires five packets and two round-trip times. Your manager suggests that this process could be streamlined if a client would begin each TCP conversation by sending a packet with the `SYN` and `FIN` flags set, plus application data in the segment body. Then the server can return one packet containing the response and the TCP connection can consist of exactly two packets.

(a) If this scheme is to be implemented, which flags should be set in the server's response to the client? Why?

(b) Explain one way in which the Berkeley Socket API makes it inconvenient to pursue this approach.

(c) As you are writing up a report for your manager, one of your co-workers drops by. Leslie claims that this approach has a fatal flaw: neither client queries nor server responses can be more than 1500 bytes long. Is this true? Explain.

Problem 5: Sized for Speed

You are hired to design a reliable byte-stream protocol that uses a sliding window (like TCP). This protocol will run over a 100 Mb/s network. The round-trip time will be bounded by 100 ms and you may assume the maximum segment lifetime is 60 seconds.

(a) How many bits wide should the AdvertisedWindow header field be? Show your work.

(b) How many bits wide should the SequenceNumber field be? Show your work.

(c) Your manager walks into your office and asks you whether your design can be made "future-proof." For each of the three design parameters, briefly explain whether or not you think it makes sense to design for a factor of 100 increase (rate = 10 Gb/s, RTT = 10s, MSL = 6000s).