Computer Network Midterm 2019

Question 1: " Basic of Computer Networks and the Internet " (15%)

a) The amount of time required to push all of a packet's bits into a link is called _Transmission time/delay_

b) The Internet's application layer includes many protocols, such as the _HTTP_ protocol (which provides for Web document request and transfer), _SMTP_ (which pro-vides for the transfer of e-mail messages), and _FTP/SFTP/P2P_ (which provides for the transfer of files between two end systems).

c) How many IP Protocols are there in the Internet? Please specify their names respectively. _IPv4/IPv6_

d) An application-layer process sends messages into, and receives messages from, the network through a software interface called a _Socket_. This is also referred to as an API which stands for _Application Programming Interface_.

e) Internet protocols are defined in _RFCs_.

f) In order to deal with the issue of scale, the DNS uses a large number of servers, organized in a _hierarchical_ fashion and distributed around the world. No single DNS server has all of the mappings for all of the hosts in the Internet. Instead, the mappings are distributed across the DNS servers. To a first approximation, there are three classes of DNS servers— _root_ DNS servers, _TLD_ DNS servers, and _authoritative_ DNS servers

g) In a DNS server, a resource record is a four-tuple that contains the following fields: (Name, Value, Type, TTL) The meaning of Name and Value depend on Type:

- If Type = A, then Name is a _hostname_ and Value is the _IP address_ for the host-name. Thus, a Type A record provides the standard hostname-to-IP address mapping. As an example, (relay1.bar.foo. com,145.37.93.126, A) is a Type A record.
- If Type = _NS_, then Name is a domain and Value is the host-name of an authoritative DNS server that knows how to obtain the IP addresses for hosts in the domain. This record is used to route DNS queries further along in the query chain.

Question 2: " End to end transmission " (10%)

How long does it take a packet of length 3,000 bytes to transmit from one end to the other end over a link of distance 2,500 km, propagation speed 2.5×10^8 m/s, and transmission rate 4 Mbps?

(Please write down your calculation process, or you will get 0 point)

$$\frac{2500 \times 10^3}{2.5 \times 10^8} + \frac{3000 \times 8}{4 \times 10^6} = 0.01 + 0.006 = 0.016(s)$$

Question 3: " Quickies " (35%)

Answer each of the following questions *briefly, i.e., in at most 3 sentences.*

a) (15%)

①(5%) What does it mean when we say that control messages are "in-band"?

Answer: It means that control message and data messages may be interleaved with each other on the same connection. Indeed, a single message may contain both control information and data

②(5%) What does it mean when we say that control messages are "out-of-band"?

Answer: It means that control and data messages are carried on separate connections

③(5%) Give an example of a protocol that has in-band control messages and one example of a protocol that has out-of-band control messages.

In-band control examples: HTTP, DNS, TCP, SMTP

Out-of-band control example: FTP

b) (5%) Consider a TCP connection between hosts A and B. Suppose that the TCP segments from A to B have source port number x and destination port number y. What are the source and destination port numbers for the segments traveling from B to A?

Answer: source port is y, destination port is x

c) (5%) What is the purpose of the connection-oriented welcoming socket, which the server uses to perform an *accept()*? Once the *accept()* is done, does the server use the welcoming socket to communicate back to the client? Explain..

Answer: a connection-oriented server waits on the welcoming socket for an incoming connection request. When that connection request arrives a new socket is created at the server for communication back to that client

d) (10%) Suppose a web server has 300 ongoing TCP connections.

①(5%) How many server-side sockets are used?

②(5%) How many server-side port numbers are used?

(Hint: remember the server implements fork() as introduced in lectures)

Answer: 301 server-side sockets; 1server-side port number

Question 4: " A reliable data transfer protocol " (25%)

Consider a scenario in which a Host A wants to simultaneously send messages to Hosts B and C. A is connected to B and C via a broadcast channel – a packet sent by A (e.g., in a single udt_send() operation) is carried by the channel to both B and C. Suppose the broadcast channel connecting A, B, and C

- can independently lose and corrupt messages from A to B and C (and so, for example, a message sent by A might be correctly received at B but not at C)
- has a maximum bounded delay of D (i.e., if a message is sent by A, it will either be lost or arrive at B and/or C within D time units).
- any control messages (e.g., an ACK or NAK) sent by B or C to A will only be received by A but can be lost or corrupted

a) (20%) Design a stop-and-wait-like error-control protocol for reliably transferring a packet from A to B and C, such that A will not get new data from the upper layer until it knows that *both* B and C have correctly received the current packet. Give a FSM description for A and B (assuming the FSM for C is similar, if it is not similar give the FSM for C as well).

Answers: This problem is a variation on the simple stop and wait protocol (rdt3.0). Because the channel may lose messages and because the sender may resend a message that one of the receivers has already received (either because of a premature timeout or because the other
receiver has yet to receive the data correctly), sequence numbers are needed. As in rdt3.0, a 0-bit sequence number will suffice here. Note that the receivers need to identify themselves in their ACK so that the sender will know which receiver sent the ACK, so that it can make sure that it has received ACKs from both receivers.

The sender and receiver FSM are shown in the figure below (note: I do not expect you to have come up with a solution at the level of syntactic detail shown below!). In this problem, the sender state indicates whether the sender has received an ACK from B (only), from C (only) or from neither C nor B. The receiver state indicates which sequence number the receiver is waiting for.

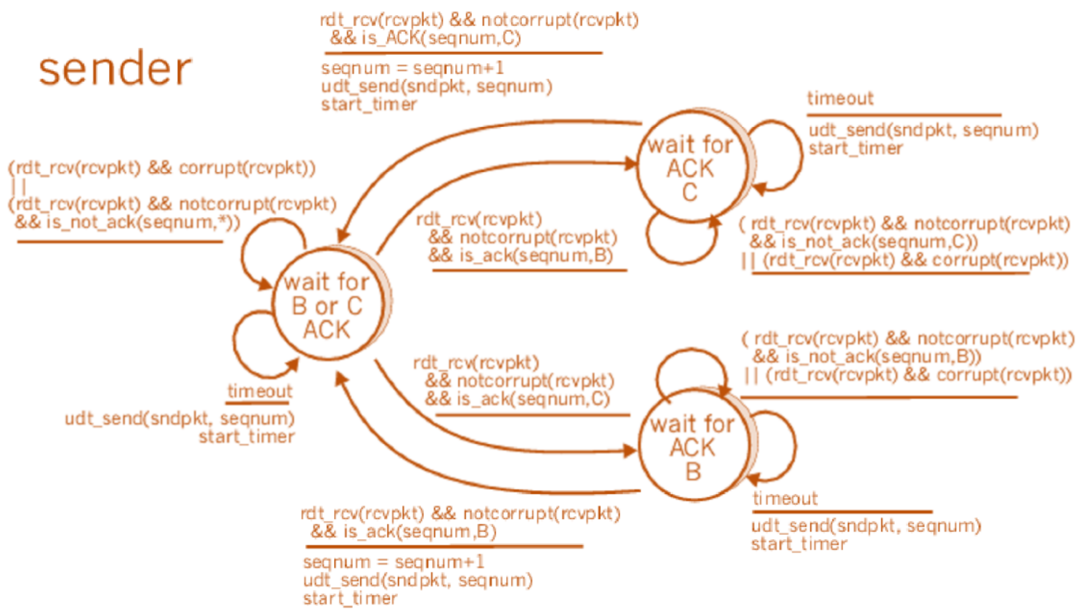

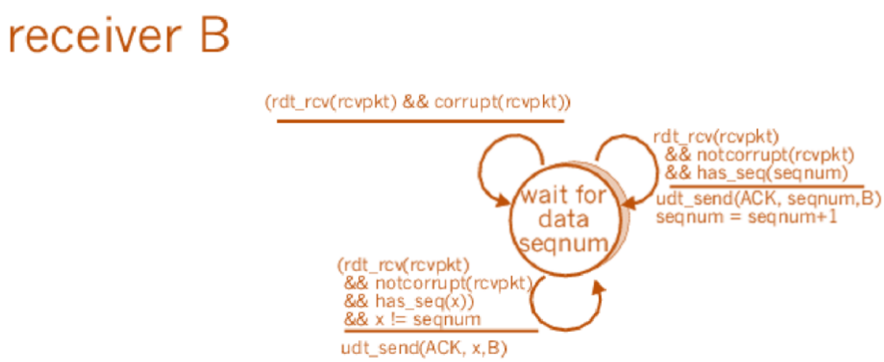

b) (5%) Also, give a description of the packet format used.

Sender-to-receiver data:

Seqnum	data

receiver-to-sender control:

acknum	B | C

Question 5: DNS service (15%)

Suppose you open a startup company "foo" and want to set up your company network. Your network has the following servers:

- DNS server: "dns1.foo.com" with IP as "140.112.30.40"
- Web server: "foo.com" with two IP as "140.112.30.55" and "140.112.30.56". The web server also has a name as "www.foo.com".

- Email server: "galaxy.foo.com" with IP as "140.112.30.60"
- Your company's email address is "username@foo.com".

a) (5%) What resource records (RRs) do you need to provide to the upper-level ".com" Registrar?

(foo.com, dns1.foo.com, NS)
(dns1.foo.com, 140.112.30.40, A)

b) (10%) What RRs do you need to put in your company's DNS server?

(foo.com, 140.112.30.55, A)
(foo.com, 140.112.30.56, A)
(www.foo.com, foo.com, CNAME)
(galaxy.foo.com, 140.112.30.60, A)
(foo.com, galaxy.foo.com, MX)

<u>Question 6: Caching and delays</u> (25%)

Consider the networks shown in the figure below. There are two user machines m1.a.com and m2.a.com in the network a.com. Suppose the user at m1.a.com types in the URL www.b.com/bigfile.htm into a browser to retrieve a 1Gbit (1000 Mbit) file from www.b.com.

a) (5%) List the sequence of DNS and HTTP messages sent/received from/by m1.a.com as well as any other messages that leave/enter the a.com network that are not directly sent/received by m1.a.com from the point that the URL is entered into the browser until the file is completely received. Indicate the source and destination of each message. You can assume that every HTTP request by m1.a.com is first directed to the HTTP cache in a.com and that the cache is initially empty, and that all DNS requests are iterated queries.

b) (5%) How much time does it take to accomplish the steps you outlined in your answer to **a)** ? Explain how you arrived at this answer. In answering this question, you can make the following assumptions

- The packets containing any DNS commands and HTTP commands such as GET are very small compared to the size of the file, and thus their transmission times (but not their propagation times) can be neglected.
- Propagation delays within the LAN are small enough to be ignored. The propagation from router R1 to router R2 is 100 ms.
- The propagation delay from anywhere in a.com to any other site in the Internet (except b.com) is 500 ms.

- *M1.a.com needs to resolve the name `www.b.com` to an IP address so it sends a DNS REQUEST message to its local DNS resolver (this takes no time given the assumptions below)*
- *Local DNS server does not have any information so it contacts a root DNS server with a REQUEST message (this take 500 ms given the assumptions below)*
- *Root DNS server returns name of DNS Top Level Domain server for .com (this takes 500 ms given the assumptions below)*
- *Local DNS server contacts .com TLD (this take 500 ms given the assumptions below)*
- *TLD .com server returns authoritative name server for b.com (this takes 500 ms given the assumptions below)*
- *Local DNS server contacts authoritative name server for b.com (this takes 100 ms given the assumptions below)*
- *Authoritative name server for b.com returns IP address of `www.b1.com`. (this takes 100 ms given the assumptions below)*
- *HTTP client sends HTTP GET message to `www.b1.com`, which it sends to the HTTP cache in the a.com network (this takes no time given the assumptions).*
- *The HTTP cache does not find the requested document in its cache, so it sends the GET request to `www.b.com`. (this takes 100 ms given the assumptions below)*
- *`www.b.com` receives the GE request. There is a 1 sec transmission delay to send the 1Gbps file from `www.b.com` to R2. If we assume that as soon as the first few bits of the file arrive at R1, that they are forwarded on the 1Mbps R2-to-R1 link, then this delay can be ignored.*
- *The 1 Gbit file (in smaller packets or in a big chunk, that's not important here) is transmitted over the 1 Mbps link between R2 and R1. This takes 1000 seconds. There is an additional 100 ms propagation delay.*
- *There is a 1 sec delay to send the 1Gbps file from R1 to the HTTP cache. If we assume that as soon as the first few bits of the file arrive at the cache, that they are forwarded to the cache, then this delay can be ignored.*
- *There is a 1 sec delay to send the 1Gbps file from the HTTP cache to m1.a.com. If we assume that as soon as the first few bits of the file arrive at the cache, that they are forwarded to the cache, then this delay can be ignored.*

The total delay is thus: .5 + .5 + .5 +.5 +.1 + .1 + 1 + 1000 +1+1 = 1105.2 secs (1002.2 is also an OK answer).

c) (5%) Now assume that machine m2.a.com makes a request to exactly the same URL that m1.a.com made. List the sequence of DNS and HTTP messages sent/received from/by m2.a.com as well as any other messages that leave/enter the a.com network that are not directly sent/received by m2.a.com from the point that the URL is entered into the browser until the file is completely received. Indicate the source and destination of each message.

[Hint: make sure you consider caching here]

d) (5%) How much time does it take to accomplish the steps that you outlined in your answer to **c)** ?

- m2.a.com needs to resolve the name www.b.com to an IP address so it sends a DNS REQUEST message to its local DNS resolver (this takes no time given the assumptions above)
- The local DNS server looks in its cache and finds the IP address for www.b.com, since m1.a.com had just requested that that name be resolved, and returns the IP address to m2.b.com. (this takes no time given the assumptions above)
- HTTP client at m2.a.com sends HTTP GET message to www.b1.com, which it sends to the HTTP cache in the a.com network (this takes no time given the assumptions).
- The HTTP cache finds the requested document in its cache, so it sends a GET request with an If-Modified-Since to www.b.com. (this takes 100 ms given the assumptions)
- www.b.com receives the GET request. The document has not changed, so www.b.com sends a short HTTP REPLY message to the HTTP cache in a.com indicating that the cached copy is valid. (this takes 100 ms given the assumptions)
- There is a 1 sec delay to send the 1Gbps file from the HTTP cache to m2.a.com.
- The total delay is thus: .1 + .1 + 1 = 1.2 secs

e) (5%) Now suppose there is no HTTP cache in network a.com. What is the maximum rate at which machines in a.com can make requests for the file www.b.com/bigfile.htm while keeping the time from when a request is made to when it is satisfied non-infinite in the long run?

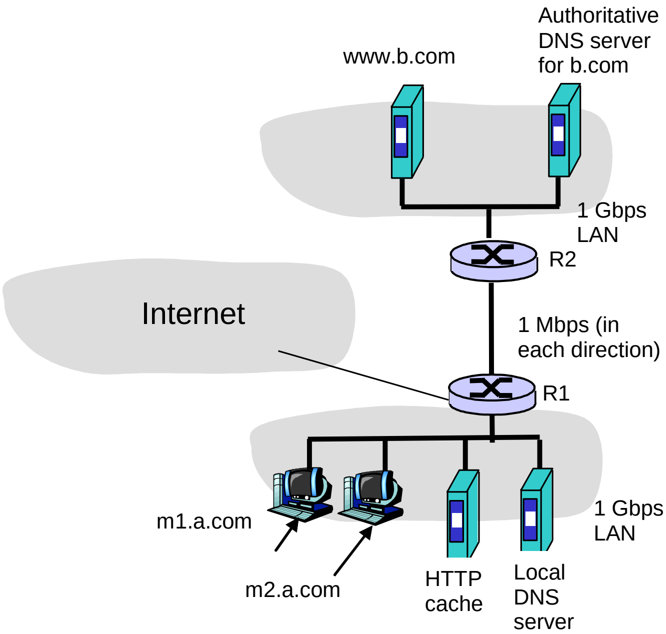

Answer: since it takes 1000 secs to send the file from R2 to R1, the maximum rate at which requests to send the file from b.com to a.com is 1 request every 1000 seconds, or an arrival rate of .001 requests/sec