Name ….....ID …..... Section ….....Seat No.....

Sirindhorn International Institute of Technology
Thammasat University

Final Examination Answers: Semester 2/2009

Course Title : ITS413 Internet Technologies and Applications

Instructor : Dr Steven Gordon

Date/Time : Wednesday 10 March 2010, 13:30-16:30

Instructions:

- This examination paper has 21 pages (including this page).
- Conditions of Examination
  Closed book
  No dictionary
  Non-programmable calculator is allowed
- Students are not allowed to be out of the exam room during examination. Going to the restroom may result in score deduction.
- Turn off all communication devices (mobile phone etc.) and leave them under your seat.
- Write your name, student ID, section, and seat number clearly on the answer sheet.
- The space on the back of each page can be used if necessary.

Questions [100 marks]

Question 1 [8 marks]

Consider the following scenario. SIIT uses voice over IP for all internal calls (that is, between SIIT staff and faculty within Bangkadi/Rangsit, as well as between Bangkadi and Rangsit). The only link between Bangkadi and Rangsit is a 2Mb/s bi-directional WiMax link (that is, 2Mb/s in one direction and 2Mb/s in the other direction). The LANs at Bangkadi and Rangsit are all 1Gb/s Ethernet – sufficient capacity for the WiMax link to be the bottleneck. If all VoIP phones are equipped to use G.726-based application (see details below), what is the maximum number of voice calls that can be supported between Bangkadi and Rangsit at the same time? You may ignore the impact of delay, jitter and packet loss, and assume the VoIP application uses RTP over UDP. Assume WiMax uses the same frame format as Ethernet. Although you should consider the frame headers, you can ignore any other overheads of Ethernet/WiMax. Clearly show all your calculations and any assumptions.

G.726 codec: Sample size: 20 Bytes; Sample interval: 5ms; RTP voice payload size: 80 Bytes

Answer

A sample interval of 5ms means 200 samples per second (200Hz). Each sample is 20 Bytes, but each RTP packet contains 80 Bytes. That is, each packet contains 4 samples. With 200 samples per second, the source must send 50 packets per second. Each packet contains: 18 bytes of Ethernet/WiMax header/trailer; 20 Bytes of IP header; 8 Bytes of UDP header; 12 Bytes of RTP header; and 80 Bytes of data. That is, 138 Bytes per packet. Therefore the source sends at a rate of 138 x 50 = 55,200 bps. With a 2Mb/s link, that allows a maximum of 36 VoIP calls at any one time.

Question 2 [17 marks]

a) A TCP source maintains a variable, called *window*, that limits the number of bytes that can be sent before having to wait for an acknowledgement. The value of *window* is calculated from two values (also referred to as windows). What are they names of the two windows? [1 mark]

Answer

Advertised Window and Congestion Window

b) How does the TCP source calculate *window* from the two windows in the answer above? [1 mark]

Answer

window = minimum(Advertised Window, Congestion Window)

c) Explain the purpose of the two windows from part (a). [3 marks]

Answer

Advertised Window is used by the receiver to limit the rate at which the sender sends so that the buffer at the receiver does not overflow (flow control).

Congestion Window is used by the sender to determine how much can be sent so that congestion doesn't occur in the network.

The following cases require you to draw a plot of the value of TCP source *window* over time. The plot does not have to be exact or to scale, but should clearly show the different phases (additive increase, slow start, responses to loss, etc.) and how they differ. That is, the shape of the plot should be clear. You should also show the slow start threshold on your plot. As a guide, the initial values of the window are given. You should assume the time on the horizontal axis is a long period, meaning the slow start threshold should be reached within that time.

d) Buffer space at receiver is infinite. No packet loss. [3 marks]

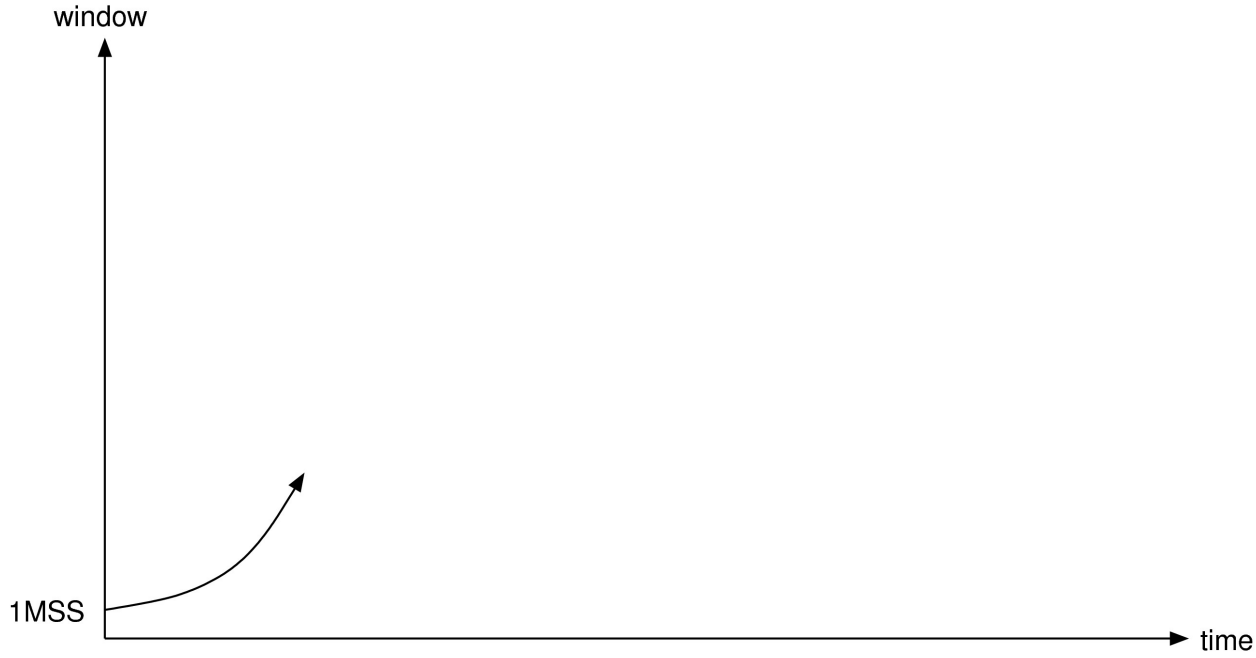

Answer

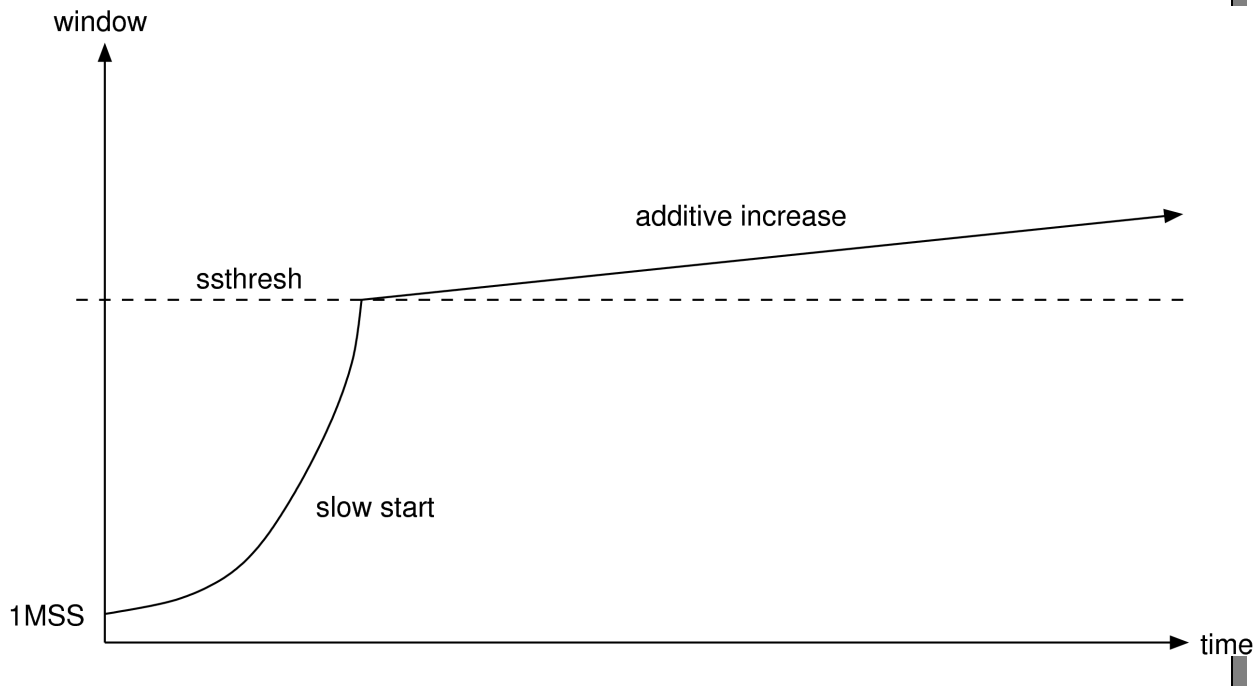

e) Buffer space at receiver is less than slow start threshold. No packet loss. [3 marks]

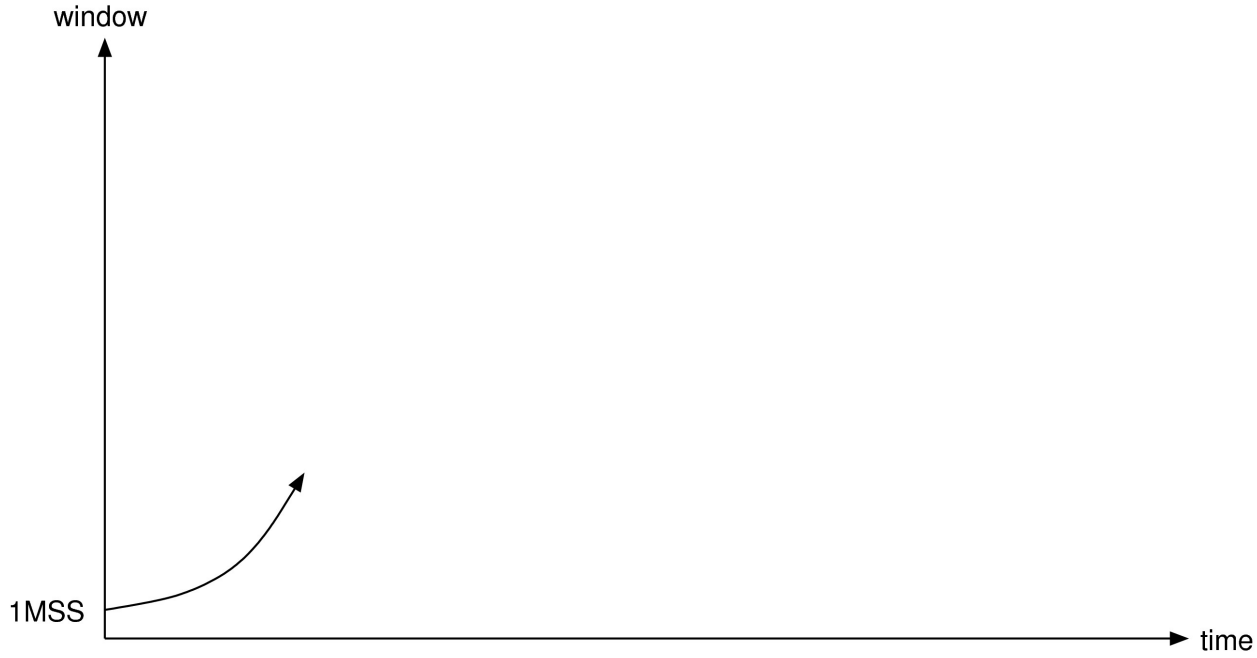

Answer

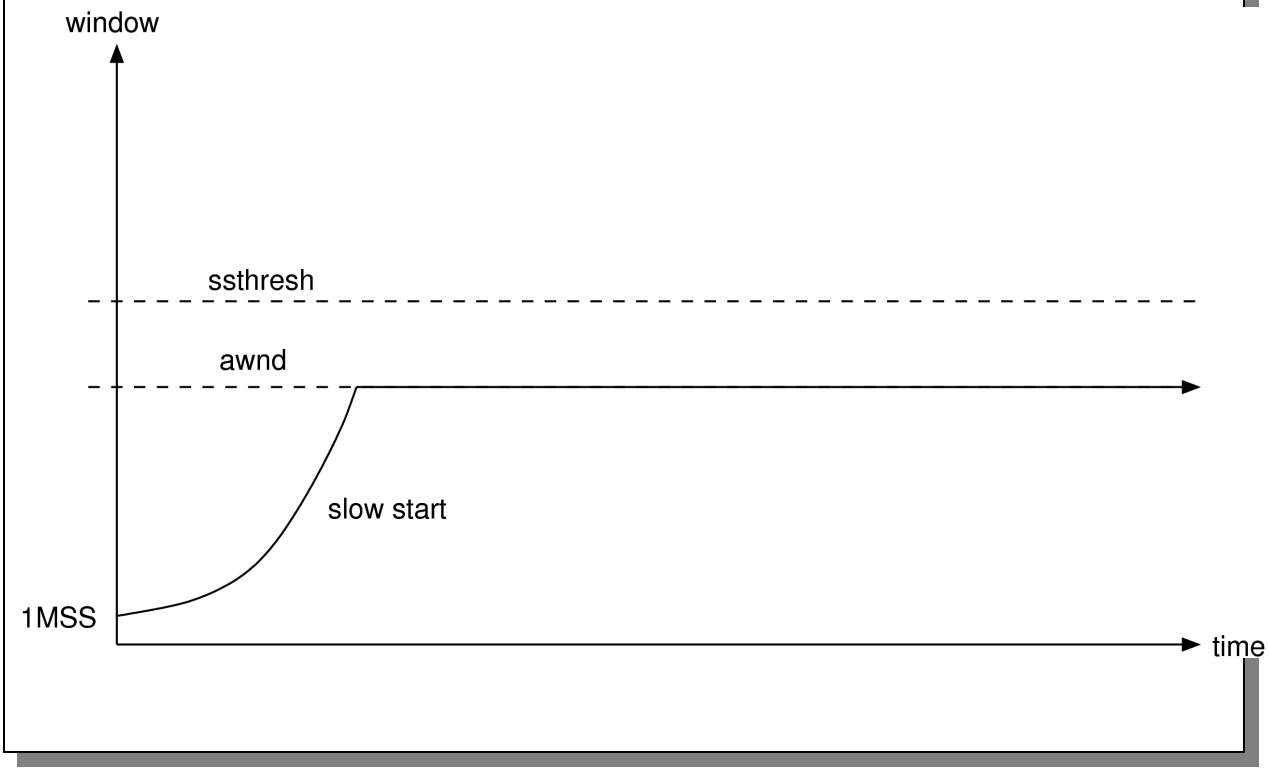

f) Buffer space at receiver is finite and greater than slow start threshold. Loss due to 3rd duplicate ACK occurs. Note that the diagram below shows the current value of the window when the loss occurs as a solid line. [3 marks]

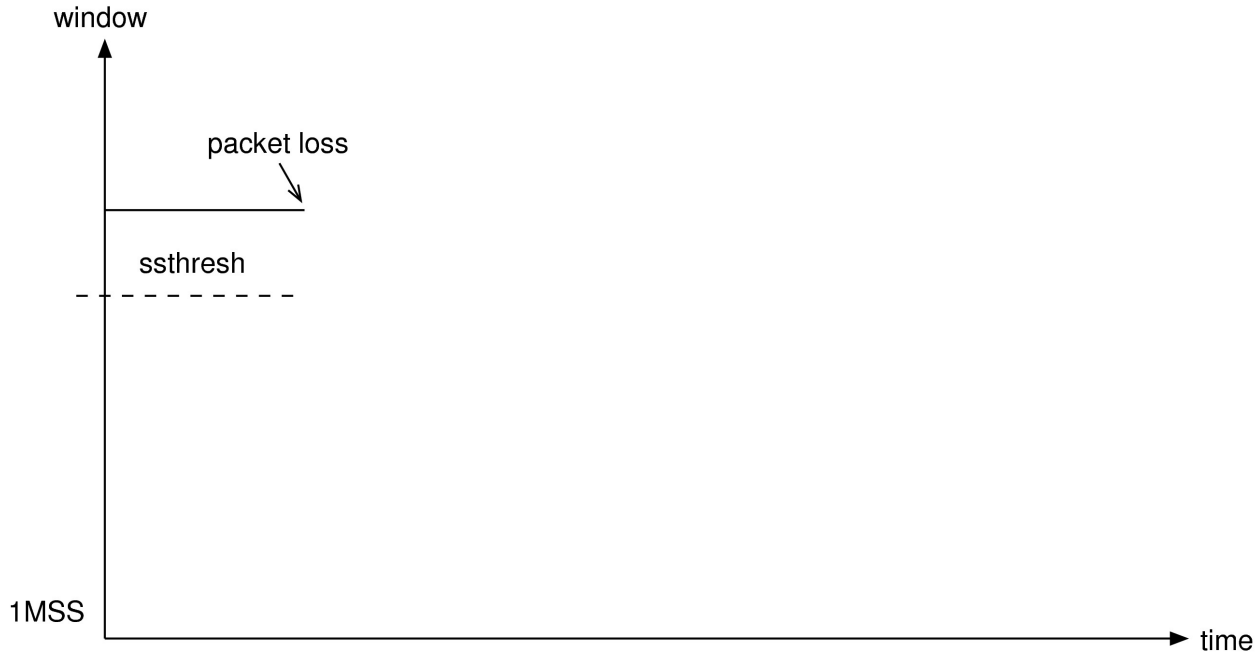

Answer

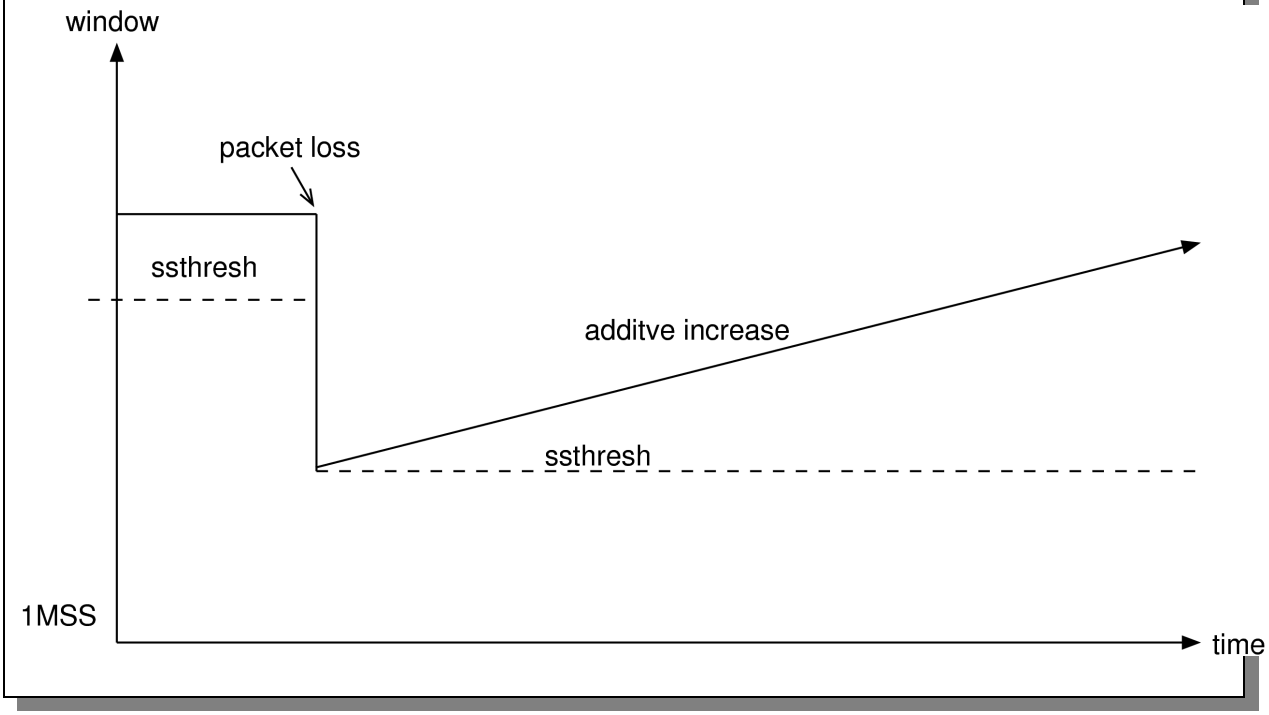

g) Explain why TCP source responds differently when there is a loss due to a timeout compared to when there is a loss due to a 3rd duplicate ACK. [3 mark]

Answer

If a single packet is lost, then the loss due to 3rd duplicate ACK should be detected before a timeout occurs. However if many packets are lost or packets are significantly delayed, then the 3rd duplicate ACK may not be received before the timeout. In other words, loss due to 3rd duplicate ACK normally means there is less congestion than if a loss due to timeout occurs. Therefore, in the event of a loss due to timeout (ie. large congestion), TCP decreases its congestion window significantly in order the avoid further congestion.

Question 3 [13 marks]

Consider the IP network in the figure below. A packet is to be sent from one source to multiple destinations.

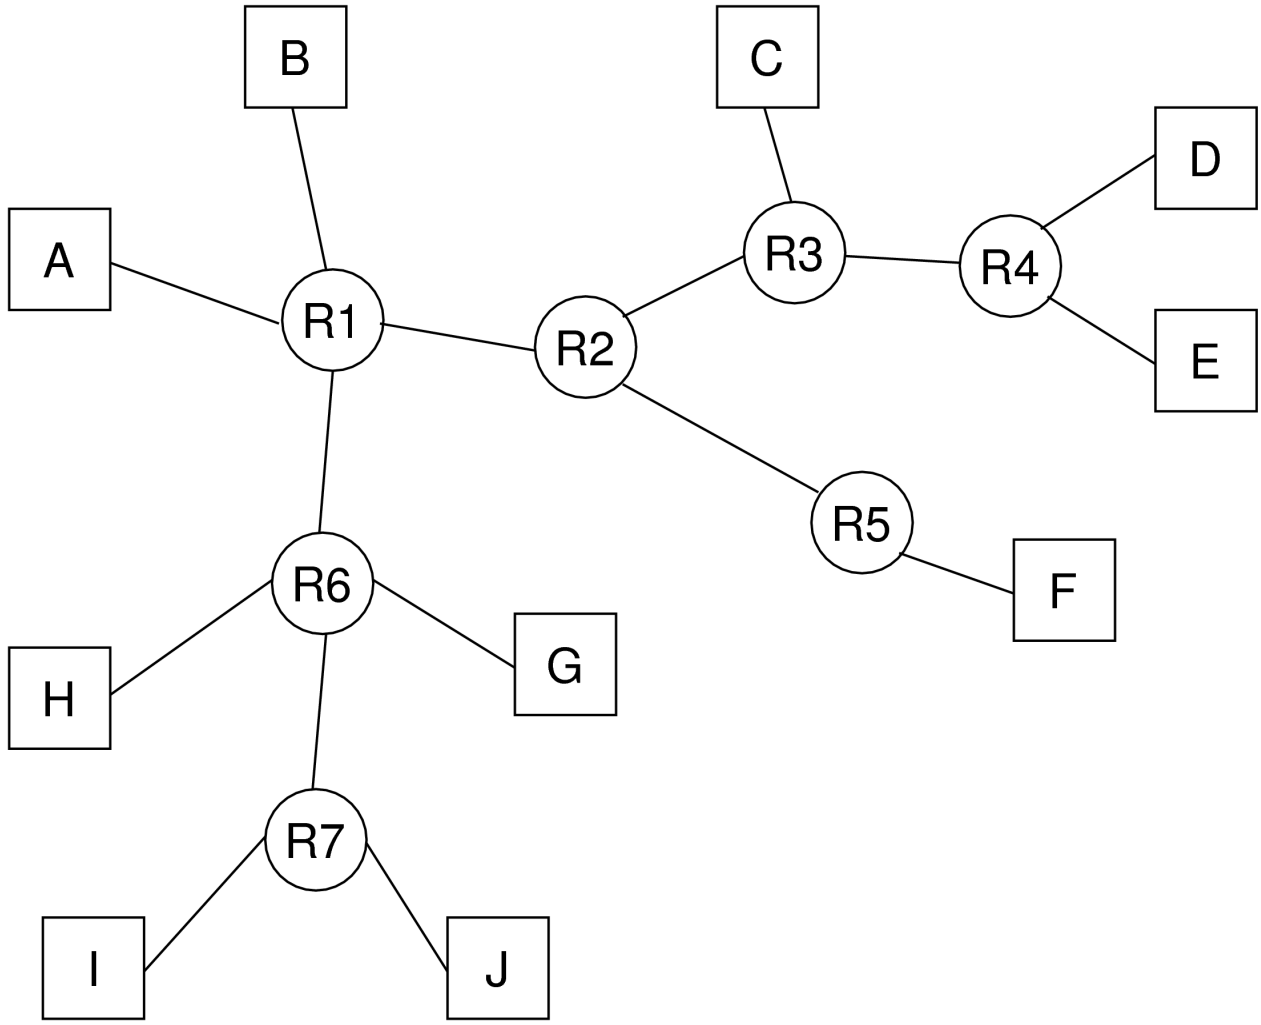

In the following cases, calculate the minimum number of packets that must be transmitted by the hosts and routers. For example, if host C sends a packet to host D, then the answer is 3: C transmits a packet to R3, R3 transmits a packet to R4 and R4 transmits a packet it D. State any assumptions if you make them.

a) Source A is to send a packet to the destination hosts B, C, D, G, J using unicast. [3 marks]

Answer

18 packets

b) Source A is to send a packet to the destination hosts B, C, D, G, J using multicast. [3 marks]

Answer

11 packets

c) Source A is to send a packet to the destination hosts B, H, G, I, J using multicast. [3 marks]

Answer

8 packets

d) If the data being sent by the source host was from a live TV event, when multicast is being used what protocol can be used to allow a new destination host (e.g. host F) to also receive the packets? [1 mark]

Answer

IGMP

e) Explain the difference between anycast, multicast and broadcast. [3 marks]

Answer

Broadcast involves sending packets to all members of a network; multicast involves sending packets to a group of members of a network (subset of all members); anycast involves sending packets to any one in a group of members of a network.

Question 4 [10 marks]

Service differentiation or soft QoS is provided by the DiffServ architecture in IP networks. Consider two DiffServ-capable hosts sending packets via a DiffServ-capable router.

a) Explain how the hosts may use *packet marking* so that, for example, the video packets from host A receive different QoS than the data packets from host B. [2 marks]

Answer

The hosts can mark packets by setting the DiffServ (or Type of Service) field in the IP header of the packets to a certain level, e.g. 5 for video, 3 for data.

b) Explain what the routers does to give the video packets of host A different QoS than the data packets of host B. [2 marks]

Answer

Once the router classifies the packet to receive a certain QoS, it may use different queueing schemes to determine when to send the packet. E.g. video packets go to the head of the queue, data packets to the end of a queue.

Hard QoS is provided by the IntServ architecture in IP networks.

c) Explain the difference between soft QoS and hard QoS. [2 marks]

Answer

Soft QoS allows priority to be given to specific packets, but cannot guarantee applications/hosts/packets will receive an absolute performance guarantee, e.g. never less than 1Mb/s. Hard QoS provides performance guarantees, e.g. an application is guaranteed to receive at least 1Mb/s.

d) Explain the concepts of *resource reservation* and *admission control* in IntServ, including how they ensure hard QoS. You should also give examples of resources reserved. [4 marks]

Answer

Resource reservation involves a host sending messages to reserve resources along the path to the destination host. The resources include portion of link capacity, buffer space and processing capability at routers. If the network has sufficient resources to satisfy the request, then the connection/application is admitted. If there are not sufficient resources then the application is rejected.

Question 5 [13 marks]

a) Explain the difference between IPTV, File-based TV distribution and Internet (or Net) TV. You should also mention the network (or type of network) that each are delivered over. [2 marks]

Answer

Net TV is the real-time delivery of low to medium quality TV over the Internet (usually via web sites). File based TV is non-real-time delivery of entire content over Internet (usually P2P file sharing). IPTV is delivery of high quality TV (equivalent to digital TV) over IP based network (usually over private IP network).

The figure below shows a typical structure of an IPTV network.

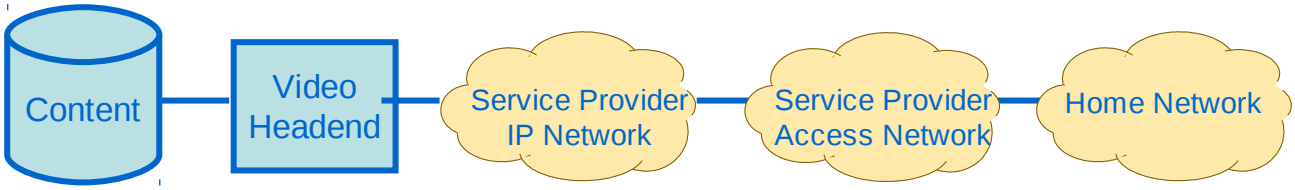

b) What portion of the network is the performance bottleneck? [1 mark]

Answer

Service Provider Access Network

c) List two network technologies that may be suitable for delivering IPTV into a home (that is, the technology used to connect your home to a larger service provider network). You must give accurate names or abbreviations of standards/protocols, referring to versions if necessary. [2 marks]

Answer

ADSL2, ADSL2+, IEEE 802.11g/a/n, Ethernet, Optical Fibre or FTTH, WiMax

d) Fibre-to-the-X may be used to refer to a range of different fibre-based technologies, where X can have different names.

i. List two possible Fibre-to-the-X technologies for IPTV [1 mark]

ii. Describe briefly the two technologies [2 marks]

iii. Explain the trade-offs (advantages and disadvantages) between the two technologies [2 marks]

Answer

Fibre-to-the-Home: deliver optical fibre directly from the Service Provider IP network to the home, bypassing the need for copper cables. Fibre-to-the-Curb: deliver optical fibre to the street level (covering 10's of homes), and then use copper to deliver from a street-level cabinet to each home; Fibre-to-the-Node: deliver optical fibre to a neighborhood level exchange (covering 100's of homes), and then use copper to deliver to each home. The trade-offs are cost of deployment (the more use of optical fibre, the high the cost, mainly for installation – hence

FTTH is most expensive, FTTN is least expensive) versus performance (copper wires only allows relatively slow speeds compared to optical fibre – hence FTTH performs the best, FTTN is the slowest).

e) For the delivery of HDTV channels to users over the Service Provider IP Network, the bandwidth requirement of the network may be measured in terms of number of channels delivered. However, for the delivery of Video on Demand to users, the bandwidth requirement of the Service Provider IP Network may be measured in terms of number of VoD subscribers. Explain why the bandwidth requirement is measured with different metrics for the different services (TV versus VoD). [3 marks]

Answer

For normal TV delivery, content is sent at scheduled times to many users, and hence multicast can be used. Therefore, approximately, each channel will require a constant bandwidth requirement across the IP network, independent of the number of users. That is, one copy of the channel will be sent across the core of the network, while multiple copies sent at the edge of the network (when delivering to individual users). For VoD, each individual user effectively has their own tailored content (own movie and their own time), and hence multicast cannot be used. A copy of each each video must be sent for each user, hence the bandwidth requirement depends on the number of users.

Question 6 [10 marks]

Two challenges of using P2P systems for sharing resources are: *searching* and *data transfer*. This question is only about searching.

a) In Napster-like P2P systems an *index* is stored on a central server. What is the index (that is, what important information does it contain)? [2 marks]

Answer

The index maps resources or keys to peers. That is, it contains a list of keys of resources and the corresponding peers that maintain the resource.

b) Where is the index information stored in a FastTrack P2P system? [2 marks]

Answer

A super-peer stores the index information for its local peers.

c) Explain how searching works in a FastTrack P2P system. As a guide, you should clearly explain the steps for all typical cases of where the queried resource may be located (e.g. the conditions when a response to a query is returned). You should state which nodes send the queries, and to what destinations and using what method. [4 marks]

Answer

A peer sends a query to its local super-peer. If the super-peer maintains an index for the resource, it immediately returns a response. If the super-peer does not have the resource in the index, then the super-peer floods the query to all other super-peers. The super-peers that maintain the index will send a response.

d) Explain an advantage that FastTrack has compared to Gnutella. ("Explain an advantage" means state what the advantage is and state why it is present in FastTrack but not Gnutella). [1 mark]

Answer

Searching can be faster because the super-peers maintain some index information and can immediately respond, whereas in Gnutella no index information is stored about other peers.

As a result of the above, another advantage is that the number of messages (overhead) is reduced.

e) Explain a disadvantage that FastTrack has compared to Gnutella. [1 mark]

Answer

In FastTrack if a super-peer fails, then its local peers cannot search. Whereas in Gnutella if one peer fails then the others can still search.

Question 7 [11 marks]

Consider BitTorrent for data transfer in P2P systems.

a) What is the difference between a seed and a leecher in a BitTorrent swarm? [1 mark]

Answer

A seed has a copy of an entire torrent, while a leecher has only a subset of the pieces of the torrent.

b) Consider three files, A, B and C, are to be shared a single torrent T. Explain the relationship between the files, torrent, blocks and pieces. [2 marks]

Answer

The files are considered as a single torrent. The torrent is divided into multiple, fixed size pieces. The pieces are divided into multiple, fixed size blocks.

c) Where can a peer find the IP address of a tracker of a swarm it wants to join to obtain torrent T? [1 mark]

Answer

In the .torrent file for the torrent T.

d) What application-level protocol is used for communications between the peer and the tracker? [1 mark]

Answer

HTTP

e) Where can a peer find the IP addresses of other peers in a swarm to obtain torrent T? [1 mark]

Answer

From the tracker (it maintains a list of peers in the swarm).

f) Assume a peer may maintain 30 TCP connections to remote peers at a time. Some of the remote peers may be *choked* while the others are *unchoked*. Explain what this means. [3 marks]

Answer

The choked remote peers cannot download data from the peer, whereas the unchoked remote peers can downloaded data from the peer.

g) Which of the following P2P systems is closest to BitTorrent: Napster, Gnutella, FastTrack?

Explain your answer. [2 marks]

Answer

Napster. For locating peers, BitTorrent still relies one or several a central server, the tracker.

Question 8 [18 marks]

Consider a P2P system using a Chord Distributed Hash Table (DHT). The possible set of peers are shown in the figure below.

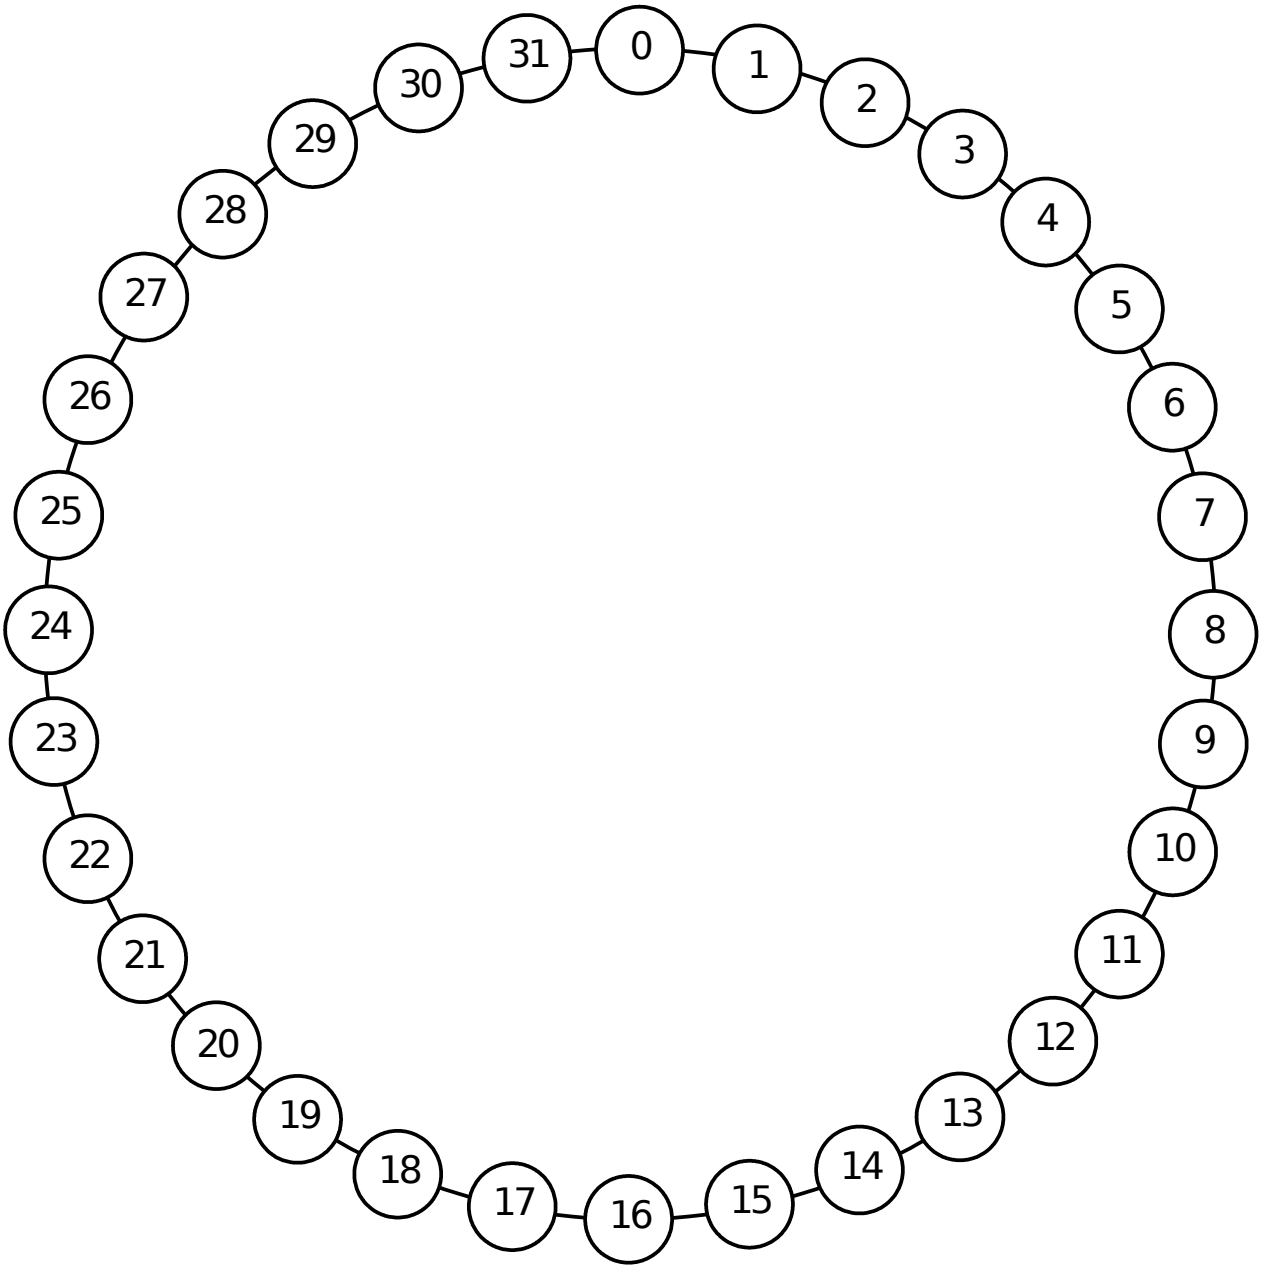

Assume currently the network contains 7 peers. The hash function H is used to determine the peer ID from the peers IP address and port number:

PeerID = H(IP + Port)

The 7 current peers have IDs: 2, 6, 8, 13, 17, 22, 30.

When peer X has a file to share, that peer calculates the hash of the filename to find the Key:

Key = H(Filename)

Peer X then sends a message to the peer that indexes that Key. The message contains the IP address and port of Peer X and the filename so the receiving peer (Peer Y) knows which peer has

the file to share. This index information is stored at Peer Y. Table 1 lists the files that have already been made available for sharing using this method.

Peer with file	Filename	Key = H(Filename)
2	siit.mpg	5
8	answers.avi	31
17	movie.avi	14
17	tu.mov	26
22	ubuntu.iso	13

Table 1: Shared Files

a) What is the maximum number of peers allowed in the network? [1 mark]

Answer

32 peers.

b) For each file, indicate the peer that maintains the index. [2.5 marks]

Filename	Indexer
siit.mpg	**6**
answers.avi	**2**
movie.avi	**17**
tu.mov	**30**
ubuntu.iso	**13**

A new peer with ID 27 joins the network. This new peer also shares a file, its413.tgz, which has a key value of 0.

c) Which peer is the indexer of the new file, its413.tgz? [0.5 mark]

Answer

Peer 2

d) Explain what else changes after peer 27 joins the network. [2 marks]

Answer

The index (key) for the file tu.mov is moved from peer 30 to peer 27.

e) Complete the Chord routing tables (that is, fill in the blank spaces). (Note that you must give the *actual* neighbour in the table). Remember, there are now 8 peers in the network. [6 marks]

Peer 2			
Positions		**Neighbour**	**Key Space**
1	3	**6**	**3**
2	4	**6**	**4..5**
4	6	**6**	**6..9**
8	10	**13**	**10..17**
16	18	**22**	**18..1**

Peer 6			
Positions		**Neighbour**	**Key Space**
1	7	**8**	**7**
2	8	**8**	**8..9**
4	10	**13**	**10..13**
8	14	**17**	**14..21**
16	22	**22**	**22..5**

Peer 8			
Positions		**Neighbour**	**Key Space**
1	9	**13**	**9**
2	10	**13**	**10..11**
4	12	**13**	**12..15**
8	16	**17**	**16..23**
16	24	**27**	**24..7**

Peer 13			
Positions		**Neighbour**	**Key Space**
1	14	**17**	**14**
2	15	**17**	**15..16**
4	17	**17**	**17..20**
8	21	**22**	**21..28**
16	29	**30**	**29..12**

Peer 17			
Positions		**Neighbour**	**Key Space**
1	18	**22**	**18**
2	19	**22**	**19..20**
4	21	**22**	**21..24**
8	25	**27**	**25..0**
16	1	**2**	**1..16**

Peer 22			
Positions		**Neighbour**	**Key Space**
1	23	**27**	**23**
2	25	**27**	**24..25**
4	26	**27**	**26..29**
8	30	**30**	**30..5**
16	6	**8**	**6..21**

Peer 27			
Positions		**Neighbour**	**Key Space**
1	28	**30**	**28**
2	29	**30**	**29..30**
4	31	**2**	**31..2**
8	3	**6**	**3..10**
16	11	**13**	**11..26**

Peer 30			
Positions		**Neighbour**	**Key Space**
1	31	**2**	**31**
2	0	**2**	**0..1**
4	2	**2**	**2..5**
8	6	**6**	**6..13**
16	14	**17**	**14..29**

f) Peer 13 is searching for the file siit.mpg. Referring to the routing tables, explain how the search is performed (that is, to which peers is the search query sent, and why). [3 marks]

Answer

The file has key 5. Peer 13 sends a search query to peer 30 since key 25 is within its key space. Peer 30 sends a search query to peer 2 since key 5 is within its key space. Peer 2 sends a search query to peer 6 since key 5 is within its key space. Peer 6 has the file and responds (either via 2 or direct to 13).

g) Explain the two important objectives of Chord using a structure of neighbours 1,2,4,8,16 ... positions away. [3 marks]

Answer

Minimise the number of neighbours to send queries to. This reduces the search time.

Minimise the number of neighbours to be maintained. This reduces the overhead of distributing routing information to discover/maintain neighbour links.

Packet Headers (for Reference)

RTP Packet

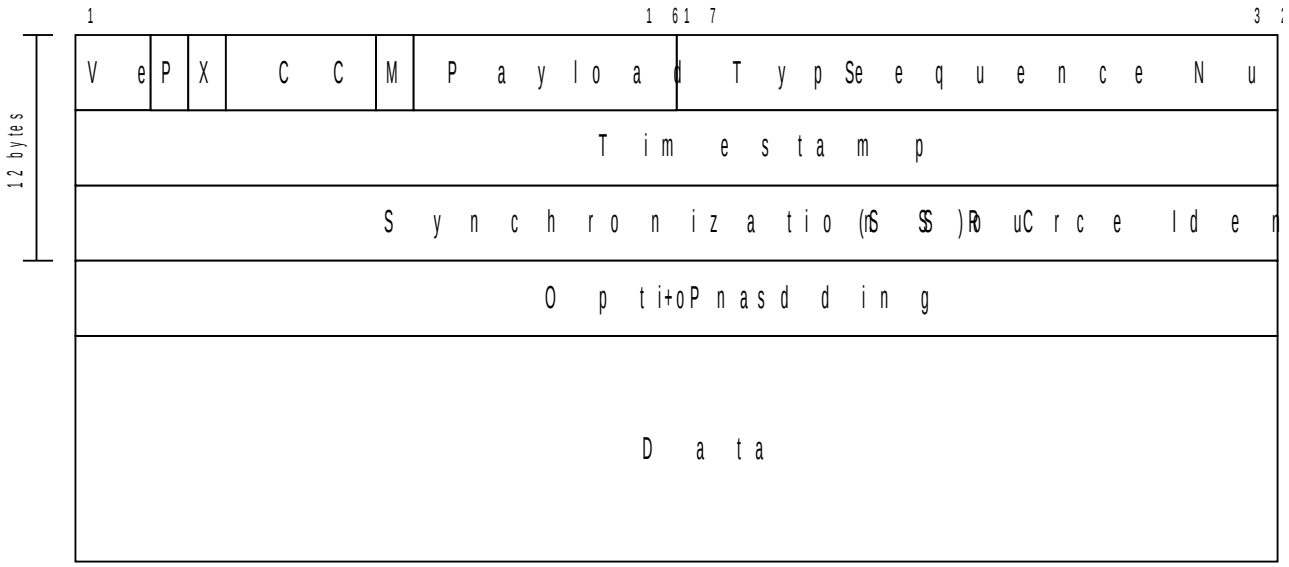

TCP Packet

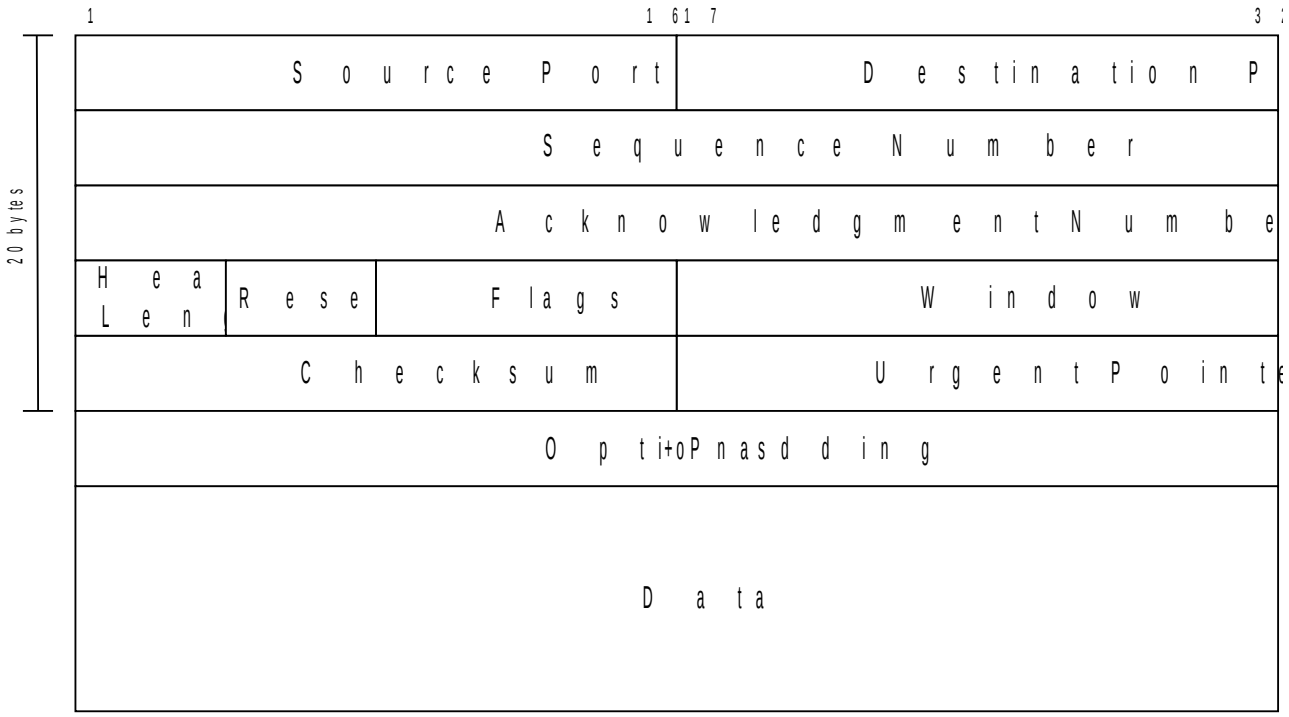

UDP Packet

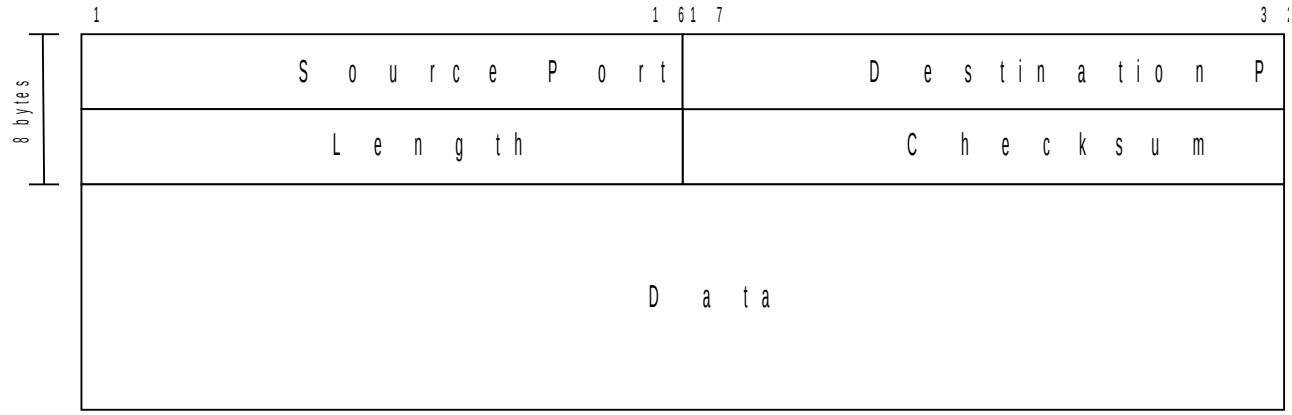

IP Packet

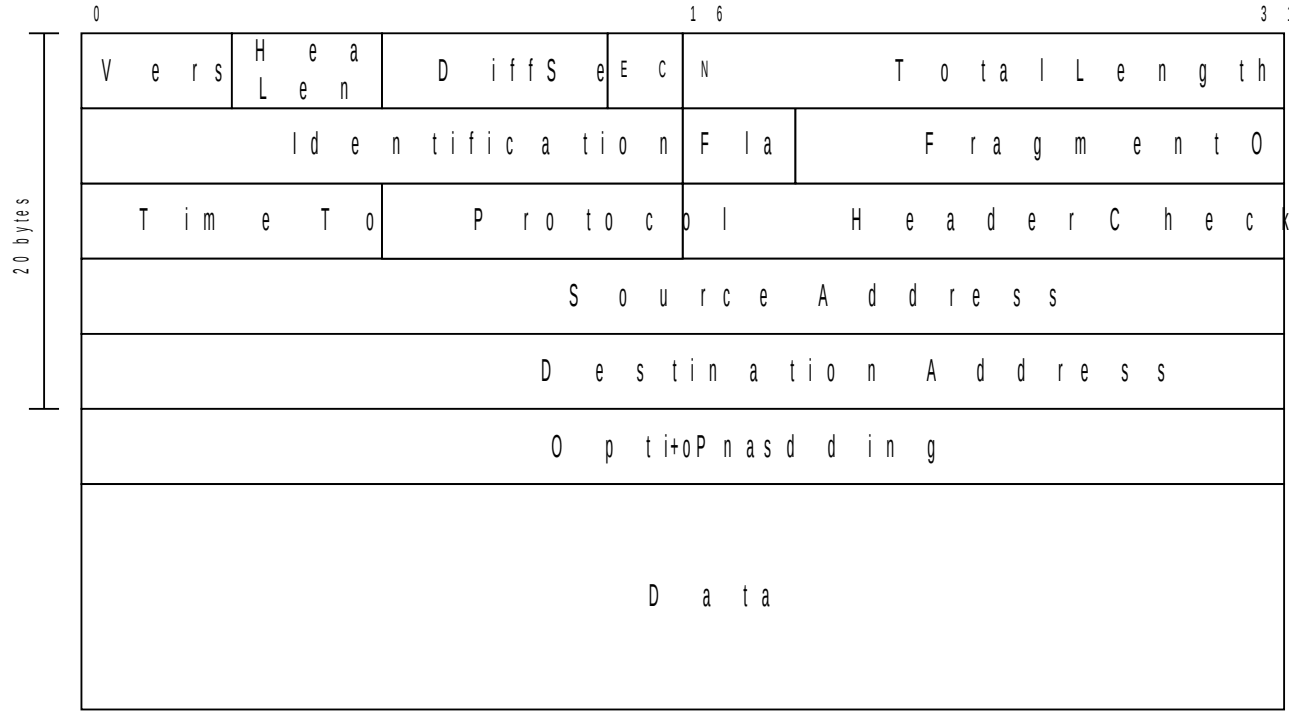

Ethernet Frame