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

Sirindhorn International Institute of Technology
Thammasat University

Final Exam Answers: Semester 1, 2012

Course Title: ITS323 Introduction to Data Communications

Instructor: Steven Gordon

Date/Time: Monday 15 October 2012; 9:00–12:00

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.
- Students are not allowed to have communication devices (e.g. mobile phone) in their possession.
- Write your name, student ID, section, and seat number clearly on the front page of the exam, and on any separate sheets (if they exist).
- Assume bits are ordered from left to right. For example, for the data 00001111, the first (1st) bit is 0 and the last (8th) bit is 1.
- Assume the speed of transmission is 3×10^8 m/s
- Reference material included at the end of the exam may be used.

Introduction to Data Communications, Semester 1, 2012

Prepared by Steven Gordon on 17 October 2012
ITS323Y12S1E02, Steve/Courses/2012/s1/its323/assessment/final-exam.tex, r2519

Question 1 [15 marks]

For each question fill in the blank space with the most appropriate term from the list below. For each blank space you must give only one answer. However, there may be more than one correct answer. You may use a term from the table in more than one question. You must not use terms that are not in the table. Each correct answer is worth 1.5 marks.

- adaptive routing
- addressing
- application
- automatic repeat request
- circuit switching
- congestion control
- datagram packet switching
- data link
- fixed routing
- flooding
- flow control
- Go-Back-N ARQ
- host
- internet
- LAN
- neighbours
- PCM
- physical
- router
- sampling period
- sampling frequency
- segmentation of data
- Selective-Reject ARQ
- sliding window
- station
- Stop-and-Wait ARQ
- subnet
- switch
- topology
- transport
- virtual circuit packet switching

(a) *Flooding* is a simple, but inefficient mechanism for sending data to all nodes in a network.

(b) If a source node sends too fast it is possible for the buffer at the destination node to become full and subsequent received packets are dropped/lost. *Flow control* is a mechanism that aims to prevent this from happening.

(c) In *Go-Back-N ARQ* if a node indicates frame i was lost, then the source retransmits frame i, as well as all frames in the window that it has previously sent after frame i.

(d) In *fixed routing* the least-cost paths are calculated when the network is designed and built; in *adaptive routing* the least-cost paths are calculated on a regular basis, e.g. every 5 minutes.

(e) An advantage of *circuit switching*, which was built for telephone networks, is that resources are reserved for the duration of the connection, meaning the application performance is guaranteed.

(f) When using *datagram packet switching*, data at the source is sub-divided into packets, and packets may take different paths to the destination.

(g) To calculate least-cost paths, a node needs to know the current network *topology*, including costs.

(h) A feature provided by TCP, but not by IP, is *flow control or congestion control.*

(i) An IP *router* is a datagram packet switch that forwards IP datagrams.

(j) A(n) *internet* is composed of multiple LANs and WANs, each using possibly different data link and physical layer technologies.

Question 2 [12 marks]

Assuming classless IP addressing is used, answer the following questions by writing your answers in the table on the next page. Unless otherwise stated, give all IP addresses in dotted decimal notation.

(a) For a host with IP address `43.109.168.14/22`: [3 marks]

  i. What is the network address?
  ii. What is the directed broadcast address?

(b) For a host with IP address `108.16.4.200` and subnet mask `255.255.0.0`: [4.5 marks]

  i. What is the network address?
  ii. What is the directed broadcast address?
  iii. What is the maximum number of IP devices that can attach to this subnet?

(c) A host does not yet have an IP address configured, nor does it know its network address. [4.5 marks]

  i. Give an IP address that the host can send to in order to send to itself.
  ii. Give an IP address that the host can send to in order to deliver an IP datagram to all nodes on its subnet.
  iii. For the case of part (ii), give the source address of the IP datagram.

Question	*Answer*
(a) i.	*43.109.168.0/22*
(a) ii.	*43.109.171.255/22*
(b) i.	*108.16.0.0/16*
(b) ii.	*108.16.255.255/16*
(b) iii.	$2^{16} - 2 = 65534$
(c) i.	*127.0.0.1*
(c) ii.	*255.255.255.255*
(c) iii.	*0.0.0.0*

Question 3 [13 marks]

Consider the network in Figure 1.

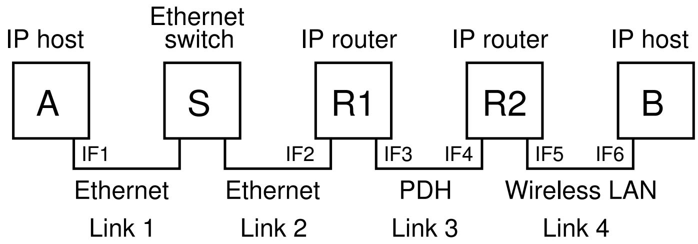

Figure 1: An internet

The addresses on each interface, IF, are given in Table 1. All subnet masks are `/24`.

Table 1: Interface Addresses

Interface	*MAC Address*	*IP Address*
IF_1	`01:23:45:67:89:AB`	`10.10.10.1`
IF_2	`AA:BB:CC:DD:EE:FF`	`10.10.10.2`
IF_3	`22:44:66:88:AA:CC`	`20.20.20.3`
IF_4	`11:33:55:77:99:BB`	`20.20.20.4`
IF_5	`FE:DC:BA:98:76:54`	`30.30.30.5`
IF_6	`00:00:11:11:22:22`	`30.30.30.6`

A user on host B is using their web browser to access a website. There are web servers running on all hosts and routers. Router R2 also acts as a DNS server. A DNS server maintains a database of mappings between domain names and IP addresses. The DNS protocol involves a client sending a DNS Query containing a domain name to the server, and the server sending a DNS Response containing the corresponding IP address back to the client. The DNS protocol uses UDP as a transport layer protocol. Assume the current values in the DNS server database are:

- `www.siit.com` → `10.10.10.1`
- `www.example.com` → `10.10.10.2`
- `www.sandilands.com` → `20.20.20.4`

The user on host B types the URL `http://www.example.com/index.html` into the address bar of their browser and presses Enter.

Assume the transport layer protocols can send the application layer data in a single segment. All requested URLs exist and are accessible on the server. Assume optional fields are not used in transport or network layer headers.

Assume you can intercept any packet on any link in the internet and see the entire contents of the packet. For the following packets and locations, complete the corresponding tables giving the value of the header fields. Some values are already given.

(a) The packet containing the DNS Query sent by host B, on link 4. [2 marks]

Layer	*Field*	*Value*
Network	Protocol	*17*
Network	Source Address	30.30.30.6
Network	Destination Address	30.30.30.5 20.20.20.4
Transport	Destination Port	*53*

(b) The 1st packet of the TCP connection setup, intercepted on link 3. [2.5 marks]

Layer	*Field*	*Value*
Network	Destination Address	10.10.10.2
Transport	Source Port	51078
Transport	Destination Port	*80*
Transport	Flags	*SYN*
Transport	Sequence Number	415

(c) The 2nd packet of the TCP connection setup, intercepted on link 3. [2.5 marks]

Layer	*Field*	*Value*
Transport	Source Port	*80*
Transport	Destination Port	*51078*
Transport	Flags	*SYN, ACK*
Transport	Sequence Number	960
Transport	Ack Number	*416*

(d) The 3rd packet of the TCP connection setup, intercepted on link 4. [1 mark]

Layer	*Field*	*Value*
Transport	Flags	*ACK*
Transport	Ack Number	*961*

(e) The packet containing the HTTP request sent by B, intercepted on link 3. [3.5 marks]

Layer	*Field*	*Value*
Data Link	Source Address	*11:33:55:77:99:BB*
Data Link	Destination Address	*22:44:66:88:AA:CC*
Network	Destination Address	*10.10.10.2*
Network	Protocol	*6*
Network	Total Length	165
Transport	Destination Port	*80*
Transport	Sequence Number	*416*

(f) The packet containing the TCP acknowledgment, acknowledging the HTTP request, intercepted on link 4. [1.5 marks]

Layer	*Field*	*Value*
Transport	Ack Number	*541*

Question 4 [12 marks]

Stop-and-wait ARQ is used on a link from A to B. A has 1,200,000 Bytes of data to send to B. Assume:

- The link propagation delay in both directions is 50 ms.
- The link data rate from A to B is 1 Mb/s.
- The link data rate from B to A is 400 kb/s.
- Frames can carry a maximum of 1200 B of data.
- Each frame has a 50 B header; an ACK does not contain any data.
- A has a timeout value of 200 ms. The timer starts immediately after the transmission of an entire data frame (i.e. after the last bit is transmitted).

(a) What is the value of the transmission delay of a DATA frame? [1 mark]

Answer. *A DATA frame contains 1200 B of data plus 50 B of header. Therefore 1250 B must be transmitted at a rate of 1 Mb/s, giving a transmission delay of 10 ms.*

(b) What is the value of the transmission delay of an ACK frame? [1 mark]

Answer. *An ACK frame contains 50 B of header and is transmitted at a rate of 400 kb/s. Hence a transmission delay of 1 ms.*

(c) Consider a single DATA frame that is lost, and then the retransmitted frame is successfully delivered. How long does it take from the start of transmission of the original DATA frame, until when the ACK is received for the retransmitted frame? [4 marks]

Answer. *The original DATA frame is transmitted, after which A starts a timer. Since the DATA frame is lost, no ACK will be returned and A will timeout 200 ms later, i.e. at time 210 ms. Then the retransmission occurs. The DATA frame is transmitted (10 ms), propagates (50 ms), ACK frame is transmitted (1 ms) and propagates (50 ms). It takes another 111 ms. Hence the total time is 321 ms.*

(d) Now consider all DATA frames. If 100 of the transmitted DATA frames are lost, what is the throughput for the data delivery? Assume no retransmitted DATA frames, nor ACK frames, are lost. That is, of all original DATA frames transmitted, 100 are lost. [4 marks]

Answer. *With 1,200,000 Bytes, there are 1000 DATA frames to be sent. For 900 DATA frames, there is no loss. The time to deliver each of this frames is: DATA transmission (10 ms), propagation (50 ms), ACK transmission (1 ms), propagation (50 ms). A total of 111 ms. For 100 DATA frames, a retransmission is needed. From the answer above, each takes 321 ms. Therefore the total time to deliver all the data is 132 seconds. The original data delivered is 1,200,000 Bytes, giving a throughput of 9,090.9 B/s or 72,727 b/s.*

(e) What is the optimal timeout value for this link? Explain your answer. [2 marks]

Answer. *The timeout value should be large enough to give enough time for the ACK to be returned if there is no packet loss, and not too large such that the source spends too much time waiting for an ACK if there is packet loss. The minimum value to allow an ACK to return is 101 ms, and hence that is the optimum value.*

Question 5 [21 marks]

Consider the network in Figure 2. The nodes/circles are packet switches.

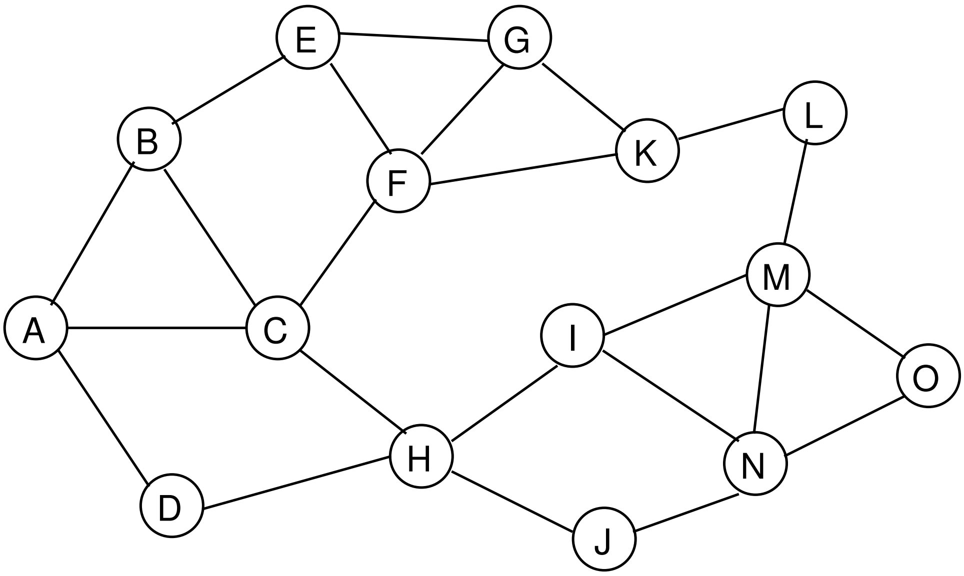

Figure 2: Network Topology

Assume flooding is to be used in the network to deliver a data packet from node A to node M.

(a) What is the optimal value of the hop limit that source node A should use? Explain your answer. [2 marks]

Answer. *The hop limit is used to limit the number of hops a packet will be forwarded. It should be large enough such that the destination will receive a copy of the packet, and small enough such that no more packets will be forwarded once the destination has received a copy. In the network the shortest path from source to destination is four hops: A—F—H—I—M. Therefore the hop limit should be four. If the hop limit was 3 or less, then the destination would not receive a copy of the packet. If the hop limit was 5 or more, although the destination would receive a copy, there would be unnecessary transmissions, e.g. from L to M and from N to M and O.*

(b) If a sequence number is used when flooding, how many copies of the packet will node H transmit? Explain your answer. Assume the delay to deliver a packet across a link is the same for all links. [2 marks]

Answer. *C and D will receive a copy from A. Then both may transmit to H. Although H receives two copies, it discards one since they have the same sequence number. H sends to I and J. H does not send back to C or D, since H knows they have already received a copy. The answer is 2.*

Assume instead of flooding, adaptive routing is to be used in the network to deliver data packets. The routing metric is not hops. The cost of a link is identical in both directions. The following least-cost paths have already been determined:

- M—L—K—G—E—B
- B—E—F—C—H—I—N—O
- J—H—C—F—E—B
- B—A—D
- G—F—C—A—D
- M—L—K—F

(c) List three (3) performance/cost metrics, other than hops, that could be used for determing the least-cost routes. [1 mark]

Answer. *Delay, throughput, capacity, error rate, security, financial cost.*

(d) What is the least-cost path from node E to node H? [1 mark]

Answer. *E—F—C—H*

(e) Draw the optimal routing tables for nodes B and F. An optimal routing table is one with the least number of rows. Each row must contain only one value for the destination and next router, with the exception that you can use * in the *destination* column to indicate "any value". In forwarding, assume the table is processed row-by-row: if one row matches, then the subsequent rows are ignored. [6 marks]

Answer. *Router B:*

- $A \rightarrow A$
- $D \rightarrow A$
- $* \rightarrow E$

Router F:

- $B \rightarrow E$
- $E \rightarrow E$
- $G \rightarrow G$
- $K \rightarrow K$
- $L \rightarrow K$
- $M \rightarrow K$
- $* \rightarrow C$

Consider virtual circuit packet switching being used to deliver packets on the path A—B—E—G—K. Assume the following:

- All packets (data, connection request, ACK) contain a header that has a transmission time of 1 ms. Data packets also contain data with an additional transmission time of 9 ms. That is, it takes a total of 10 ms to transmit a data packet.
- Link propagation delay is 2 ms in each direction.
- All links are the same, i.e. same propagation and transmission delays.
- There are 1000 data packets to be sent from A to K.
- Node A initiates the connection/circuit at time 0 seconds.

(f) At what times does node A know the connection is setup (and can start transferring data)? [2 marks]

Answer. *Connection setup involves a request packet being sent from A to K (4 hops) and a response being returned from K to A (4 hops). Delivery of a packet across a hop takes 3 ms (transmission plus propagation), and therefore the total time is 24 ms.*

(g) At what time has K received all the data? [4 marks]

Answer. *Node A transmits 1000 packets, taking 10,000 ms. The propagation of the last packet to B takes an additional 2 ms. B transmits packets to E, except most of the transmission is happening at the same time B is receiving packets from A. Its only the last packet that contributes to the total delay, i.e. an extra 10 ms transmission and 2 ms propagation. Similar for E to G and G to K. A total time of* $(10000 + 2 + 10 + 2 + 10 + 2 + 10 + 2)$ *or 10038 ms. Also considering the initial 24 ms for connection setup, K has received all the data at time 10,062 ms.*

Consider datagram packet switching being used to deliver packets on the path A—B—E—G—K. All packets must contain a 20 Byte header. They may contain a varying amount of data. The source A has 4,000 Bytes of data to send to K.

(h) Explain three (3) factors that are important in node A selecting the amount of data that should be carried in each packet. [3 marks]

Answer.

- *The larger the data the lower the overhead of the data. E.g. a packet containing 4,000 Bytes of data requires just one header. Two packets containing 2,000 Bytes of data each requires two headers; more overhead.*
- *The smaller the data the more packets can be potentially transmitted in parallel. In this case there are four links. If there are at least four packets (e.g. 1,000 Bytes per packet), then its possibly that a packet is being transmitted by each node in the path at the same time. G is transmitting packet 1, E is transmitting packet 2, B is transmitting packet 3, and A is transmitting packet 4.*

- *The smaller the data the less needs to be retransmitted in case of error.*
- *The smaller the data the more likely the receiver will have space in its buffer to store that packet.*

Question 6 [8 marks]

(a) What organisation is responsible for developing and maintaining the standards for the most popular wireless LAN technology? [1 mark]

Answer. *IEEE*

(b) Draw the protocol stack, labelling the layers with specific protocols (not just names of layers) of a laptop accessing a web site using wireless LAN. [2 marks]

Answer. *Layers, from top to bottom: HTTP; TCP; IP; IEEE 802.11 MAC; IEEE 802.11 Phy*

CSMA is used in the MAC layer of wireless LAN. The basic procedure can be described as follows:

- A station senses if anyone is transmitting. If no-one else is transmitting for a period of *DIFS* then move to next step. Otherwise, wait until they stop transmitting and restart from first step.
- Choose a random integer, R, between 0 and 15 (inclusive), and wait for $R \times SlotTime$ (this is called *Backoff*). If no-one else is transmitting during this time then move to next step. Otherwise, wait until they stop transmitting and restart from first step.
- Transmit DATA frame.
- Upon complete reception of a DATA frame, the receiver waits for a period of *SIFS* and then transmits an ACK frame.

These steps, in the case that no-one else is transmitting, are summarised in Figure 3.

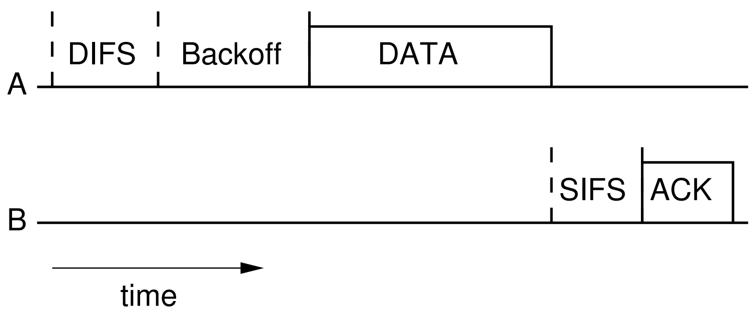

Figure 3: Basic operation of CSMA

A common set of values used for the parameters is given in Table 2.

(c) What is the best-case throughput, with one station transmitting to one other station, that can be achieved using CSMA? [5 marks]

Table 2: CSMA Parameter Values

Parameter	*Value*
Data rate	54 Mb/s
DIFS	28 μs
SIFS	10 μs
SlotTime	9 μs
DATA frame header	34 Bytes
DATA frame maximum payload (data)	1500 Bytes
ACK frame	14 Bytes

Answer. *The best case is when there is no interference, no-one else transmitting. In that case, to deliver 1500 Bytes of data in a DATA frame the time is: DIFS + Backoff + DATATransmission + SIFS + ACKTransmission. The DATA transmission time is:*

$$DATA = \frac{(34 + 1500) \times 8}{54,000,000}$$

which is 227.26 μs. The ACK transmission time is:

$$ACK = \frac{14 \times 8}{54,000,000}$$

which is 2.07 μs.

The Backoff depends on the random integer chosen. The minimum is 0 and maximum is 15. On average the value will be 7.5. Lets consider the average value. The total time is:

$$28 + 7.5 \times 9 + 227.26 + 10 + 2.07 = 334.83\mu s$$

Hence the throughput is 35.84 Mb/s. If a single packet was considered, the minimum time ($R = 0$) would be 267.33 μs, giving a throughput of 44.89 Mb/s.

Question 7 [7 marks]

Analog data, with amplitude ranging from 0 to some maximum amplitude, is to be encoded using PCM. A portion of the analog data for the first X seconds is shown in Figure 4.

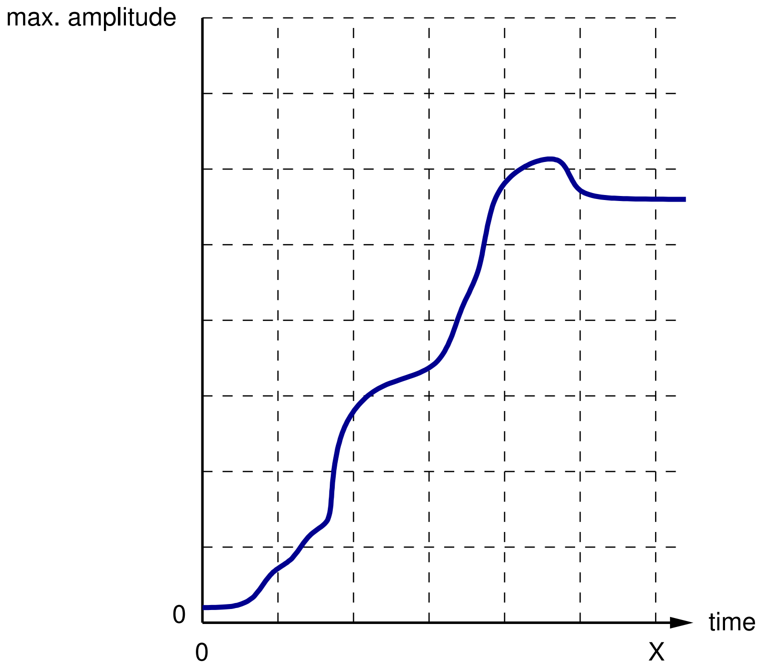

Figure 4: Analog data

Assume $X = 6\mu s$, and a sampling frequency of 500 kHz with 3-bit samples is used.

(a) What is the data transmitted? Consider only the analog input from time 0 to X, inclusive. [3 marks]

Answer. *Since $X = 6$ μs, on the x-axis each dashed line is at intervals of 1 μs. With a sampling frequency of 500 kHz, the sample period is 2 μs. That is, there are 4 samples taken, at times: 0, 2, 4, 6.*

With 3-bit samples, there are 8 possible code points. Therefore on the y-axis each dashed line is at intervals of 1.

Now at each sample point (starting at time 0), the input value is mapped to a corresponding code point.

At time $t = 0$, code point is 0, therefore the bits are 000

$t = 2$, code point $= 2$, bits $= 010$

$t = 4$, code point $= 5$, bits $= 101$

$t = 6$, code point $= 5$, bits $= 101$

The correct answer is: 000010101101.

(b) At what data rate should the data be sent so that the receiver can receive and re-construct the analog data with the same timing? [2 marks]

Answer. *3 bits are generated every sample period, i.e. every 2 μs. That is a data rate of 1.5 Mb/s.*

(c) Give an advantage and disadvantage of increasing the sampling frequency to 1 MHz: [2 marks]

i. Advantage:

Answer. *The quality of the received data will increase*

ii. Disadvantage:

Answer. *The data required to deliver the data will increase*

Question 8 [6 marks]

A sliding window protocol is used over an error free link, with a window size of 6 data frames. Each data frame has a transmission delay of 110 us, and the header is 10% of the total size (the remaining 90% of the data frame is payload, i.e. real user data). ACK transmission time is 5 us, while the link propagation delay is 290 us.

Assuming the source always has data to send, what is the maximum efficiency, expressed as a percentage, that can be achieved in delivering user data across the link?

Answer. *With sliding window flow control, the source can send 6 frames, after which it must wait for the ACK of the first frame before it can send another frame. The arrival time of the first ACK is important. If it arrives after the 6th frame has been sent, then the source must spend some time waiting for the ACK. But if it arrives during the transmission of the 6 frames, then after transmitting the 6th frame, the source can immediately transmit another frame. Therefore lets calculate which of the two cases it is.*

The time to transmit 6 frames is $6 \times 110 = 660$ *us. So if the source starts transmitting at time 0, it will finish transmitting the 6 frames at time 660.*

The time of arrival of the first ACK is: $110 + 290 + 5 + 290 = 695$ *us. That is, the source is allowed to send another frame (in the next window) at time 695.*

In this case the source sends 6 frames and then must spend some time waiting (i.e. not sending) before it is allowed to send another frame.

For every 695 us, the source is transmitting for 660 us and waiting for the rest of the time. That is, 94.96 efficient in sending frames. But note that each frame contains 90% of user data (the rest is header), and so the overall efficiency is 85.46.

Question 9 [6 marks]

TCP includes two error control mechanisms: basic retransmit and fast retransmit.

(a) What ARQ scheme is basic retransmit most similar to? [1 mark]

Answer. *Selective-reject (or selective-retransmission) ARQ, since a source is allowed to send a window size of frames (hence not stop-and-wait ARQ) and if a segment is lost only the lost segment needs to be retransmitted—other segments are buffered (hence not go-back-N ARQ).*

(b) Explain how basic retransmit works in TCP. [2 marks]

Answer. *When the source sends a segment it starts a timer. If an ACK is received before the timer expires than the timer is stopped, and the segment successfully delivered. If the timer expires before an ACK is received then the source retransmits the segment.*

(c) Explain how fast retransmit works in TCP. [2 marks]

Answer. *If the source receives four ACKs with the same acknowledgement number, then it assumes the segment that contains that acknowledgement number is lost and hence retransmits the segment.*

(d) Explain an advantage of fast retransmit (compared to basic retransmit). [1 mark]

Answer. *The source potentially retransmits a lost segment earlier than if waiting for a timeout, and hence spends less time waiting and leads to higher efficiency.*

Reference Material

Selected well-known ports:

- FTP 20 and 21
- SSH 22
- Telnet 23
- SMTP 25
- DNS 53
- HTTP 80
- HTTPS 443

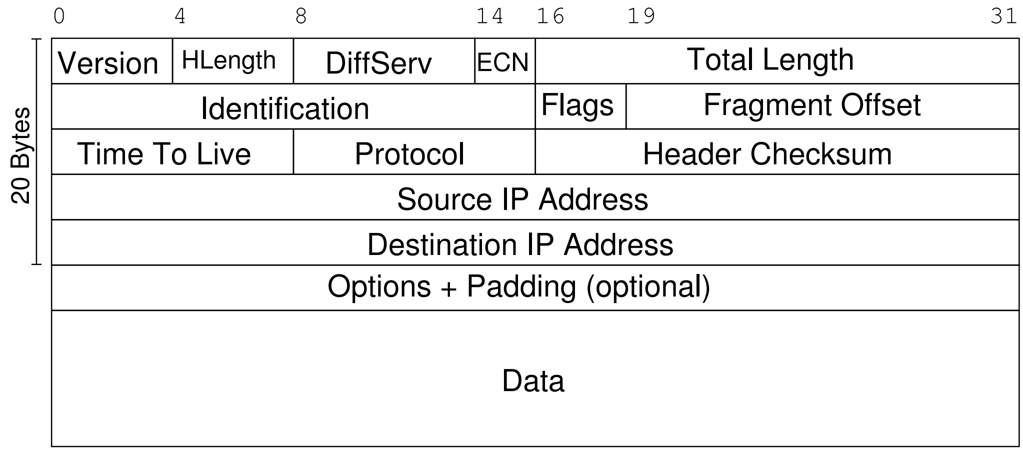

Figure 5: IP Datagram Format. Flags: Reserved, Don't Fragment, More Fragments

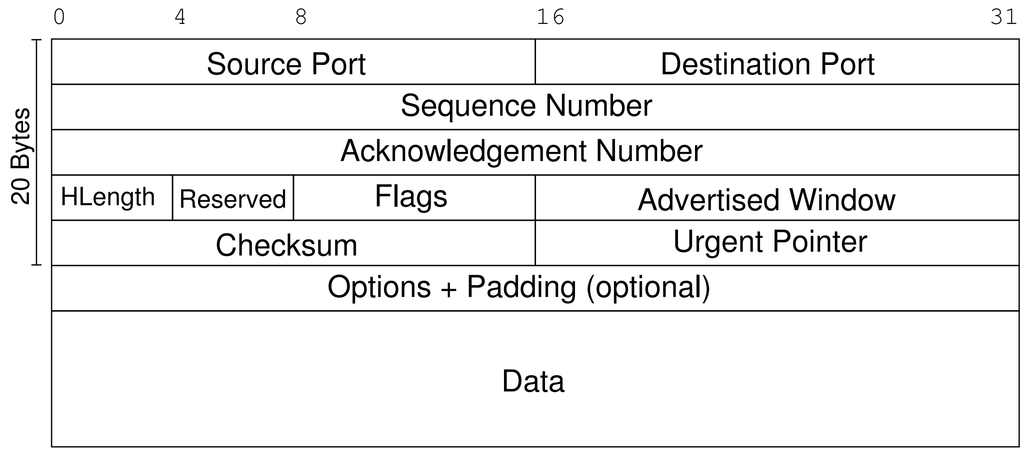

Figure 6: TCP Segment Format. Flags: CWR, ECE, URG, ACK, PSH, RST, SYN, FIN

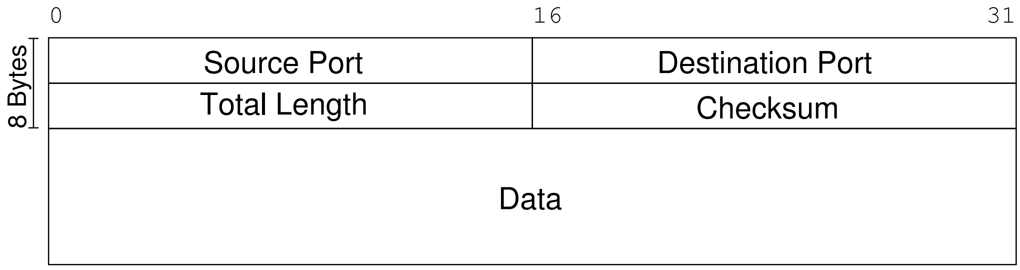

Figure 7: UDP Datagram Format

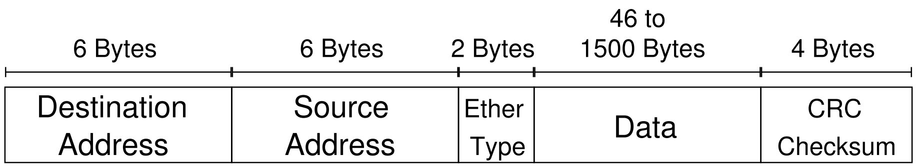

Figure 8: Ethernet Frame Format

Selected Protocol numbers:

- 1 ICMP
- 6 TCP
- 17 UDP

Selected HTTP Status Codes:

- 200 Ok
- 304 Not Modified
- 401 Unauthorized
- 404 Not Found