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

Sirindhorn International Institute of Technology
Thammasat University

Final Exam Answers: Semester 1, 2014

Course Title: ITS323 Introduction to Data Communications

Instructor: Steven Gordon

Date/Time: Monday 15 December 2014; 13:30–16:30

Instructions:

- This examination paper has 18 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 at the front of the examination room.
- The examination paper is not allowed to be taken out of the examination room. A violation may result in score deduction.
- 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, 2014

Prepared by Steven Gordon on 15 December 2014
ITS323Y14S1E02, Steve/Courses/2014/s1/its323/assessment/final-exam.tex, r3461

Question 1 [16 marks]

For each question fill in the blank space with the most appropriate term from the table 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 mark.

1 Gb/s	48-bit	HTTP	ring
1 Mb/s	64-bit	IP	routers
10 Gb/s	access points	laptops	SMTP
10 Mb/s	application	mesh	subnets
10 Mb/s	application layer	modems	TCP
100 Mb/s	bus	network layer	TDM
12-bit	circuit	NIC	transport layer
128-bit	coaxial cable	optical fibre	twisted pair
16-bit	data link layer	OS	UDP
24-bit	datagram packet	PCs	virtual circuit packet
32-bit	FDM	physical layer	-

SG: There were two answers missing from the above table: IEEE and star. Hence parts (a) and (c) were marked correct for everyone, no matter what answer they gave.

Ethernet is the common name for wired LANs. The standards for Ethernet are maintained by the standards organisation *IEEE*. The Ethernet standards focus on two layers: these two layers are normally implemented in the *NIC*.

A *star* topology is most commonly used for Ethernet LANs, where each station has a link to an Ethernet switch. Most PCs and laptops sold today have an Ethernet NIC that supports a maximum data rate of *1 Gb/s*. Most links used in Ethernet are point-to-point and use *twisted pair* as the transmission media.

There are different types of WAN technologies, including PDH, SDH, ATM and WiMax. Many WAN point-to-point links use multiplexing, which has two types. One type of multiplexing, called *FDM*, requires the bandwidth of the WAN link to be greater than the sum of the bandwidths required by the users. This type of multiplexing was widely used in original telephone networks, which used *circuit* switching. The other type of multiplexing is called *TDM*.

An internet is made up of multiple *subnets* connected together via *routers*. IP is an common internetworking protocol; it uses *datagram packet* switching. IP is normally implemented in the *OS*.

IP does not provide reliable data transfer. Instead, applications that require reliability usually use the protocol called *TCP*, which is normally implemented in the *OS*. The layers which only need to implemented on end hosts (and not necessarily on intermediate nodes) are: *transport layer* and *application layer*.

Question 2 [19 marks]

Consider the network topology in Figure 1. It contains four switching nodes, RA, RB, RC and RD, each connected via point-to-point WAN links. The link characteristics of data rate and delay are shown next to each link; the characteristics as the same in each direction.

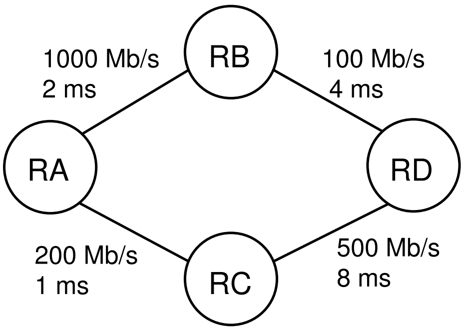

Figure 1: Network Topology 1

(a) What is the least cost path from RA to RD (and its cost) if the metric is a function of data rate defined as $cost = \frac{2000}{data\ rate}$? [2 marks]

Path: *RA–RC–RD* Cost: *14*

(b) Draw the distributed routing tables for each node if the metric is delay. That is, draw four routing tables, one for RA, another for RB and so on, where the destinations are the other nodes. [4 marks]

Answer. *Routing tables are:*

RA

Dest	*Next*
B	B
C	C
D	B

RB

Dest	*Next*
A	A
C	A
D	D

Now consider the network topology in Figure 2. The four switching nodes are IP routers. In addition there are hosts connected via switched Ethernet LANs: the subnet to the left of RA uses 1 Gb/s and the subnet to the right of RD uses 100 Mb/s. Assume the point-to-point WAN links between the routers are 10 Gb/s Ethernet.

RC

Dest	*Next*
A	A
B	A
D	A

RD

Dest	*Next*
A	B
B	B
C	B

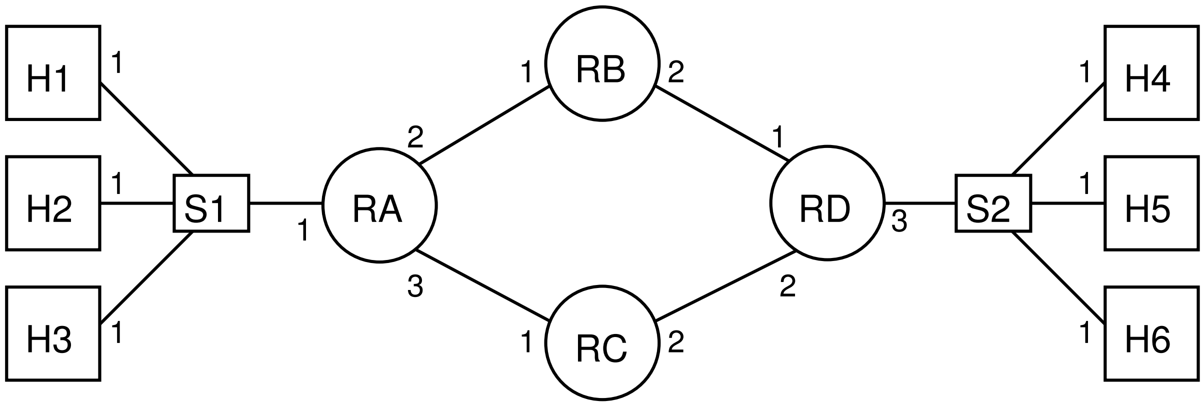

Figure 2: Network Topology 2

The numbers next to each link into a device in Figure 2 are interface numbers. For example, RA has three interfaces: interface 1 connects to the switch, interface 2 connects to RB, and interface 3 connects to RC. Table 1 lists the MAC and IP addresses currently assigned to a selection of the interfaces.

Table 1: Interface Addresses

Device	*Interface*	*MAC*	*IP*
H1	1	`11:22:33:aa:bb:cc`	`103.17.48.97/24`
H2	1	`22:33:44:bb:cc:dd`	`103.17.48.12/24`
H3	1	`33:44:55:cc:dd:ee`	?
RA	1	`12:34:56:78:90:ab`	`103.17.48.11/24`
RA	2	`cd:ef:12:34:56:78`	`24.114.6.49/16`
RA	3	`90:ab:cd:ef:12:34`	`206.100.16.1/16`
RB	1	`aa:bb:cc:11:22:33`	`24.114.120.1/16`
RB	2	`bb:cc:dd:22:33:44`	`150.12.67.5/20`
RC	1	`cc:dd:ee:33:44:55`	`206.100.16.2/16`
RC	2	`dd:ee:ff:44:55:66`	`96.27.1.1/18`
RD	1	`56:78:90:ab:cd:ef`	`150.12.67.6/20`
RD	2	`ef:cd:ab:90:78:56`	`96.27.1.2/18`
RD	3	`34:12:ef:cd:ab:90`	`97.33.180.1/21`
H4	1	`99:a1:b2:c3:d4:e5`	`97.33.177.1/21`
H5	1	`88:f6:a7:b8:c9:d0`	`97.33.179.45/21`
H6	1	`77:e1:f2:a3:b4:c5`	`97.33.181.65/21`

The following sub-questions give a scenario of one device sending data to others. For that given scenario, you need to fill in the details of the packet (Ethernet frame and/or IP datagram), that is, the packet header field values and other information about the packet. Assume the default/typical headers are used (see the Reference Material).

(c) H1 has TCP data to send to H4. The entire TCP segment (including header) is 1000 Bytes long. The path that the IP datagram takes is H1→RA→RB→RD→H4 (this path, may or may not be a least cost path identified in parts (a) or (b)). The initial time to live set by H1 is 10. Fill in Table 2 for the packet sent by H1. [4 marks]

(d) The same as for part (c), but fill in Table 3 for the packet sent by RA. [3 marks]

Answer. *See table.*

(e) H3 has just booted and currently does not have an IP address, nor does it know the network address of its current subnet or the IP address of any devices on its subnet. H3 sends an IP datagram to everyone on its subnet, with the hope that one device will respond giving H3 an IP address. Fill in Table 4 for the packet sent by H3. [2 marks]

(f) H1 wants to send one IP datagram and have it delivered to all devices on the same subnet as H4. Fill in Table 5 for the packet sent by H1. [2 marks]

Table 2: Answer for part (c)

Field	*Value*
Ethernet Source	*11:22:33:aa:bb:cc*
Ethernet Destination	*12:34:56:78:90:ab*
IP Source	*103.17.48.97*
IP Destination	*97.33.177.1*
IP Protocol	*6*
IP Total Length	*1020*

Table 3: Answer for part (d)

Field	*Value*
Ethernet Source	*cd:ef:12:34:56:78*
Ethernet Destination	*aa:bb:cc:11:22:33*
IP Source	*103.17.48.97*
IP Destination	*97.33.177.1*
IP TTL	*9*

Table 4: Answer for part (e)

Field	*Value*
IP Source	*0.0.0.0*
IP Destination	*255.255.255.255*

Table 5: Answer for part (f)

Field	*Value*
IP Source	*103.17.48.97*
IP Destination	*97.33.183.255*

Answer the following questions for the network topology in Figure 2 and Table 1.

(g) What is the network address for the subnet that H1 is attached to? [1 mark]

Answer. *103.17.48.0/24*

(h) Currently there are 4 devices on the subnet that H1 is attached to. How many *more* new devices can be added to the subnet? [1 mark]

Answer. *250*

Question 3 [6 marks]

Assume a 4-bit sequence number is used in a sliding window flow control protocol (that is, maximum window size is 15). The current state of a source node is:

- Last frame ACKed = 5
- Current window size = 10 frames

Then the node transmits 6 DATA frames, and then receives an ACK (Receive Ready) frame with number 9. After these frames have been transmitted/received, what is the new value of:

(a) Last frame ACKed: *8*

(b) Last frame transmitted: *0*

(c) Current window size: *7*

Answer. *The current state (to the left of | is last frame ACKed; in between [and] is the current window):*

0 1 2 3 4 5 | 6 7 8 9 10 [11 12 13 14 15 0 1 2 3 4] 5 6 7 8 9 10

After 6 DATA frames are transmitted it becomes:

0 1 2 3 4 5 | 6 7 8 9 10 11 12 13 14 15 0 [1 2 3 4] 5 6 7 8 9 10

After receiving an ACK with number 9 it becomes:

0 1 2 3 4 5 6 7 8 | 9 10 11 12 13 14 15 0 [1 2 3 4 5 6 7] 8 9 10

Question 4 [14 marks]

You have the task of designing a lecture recording system, where the audio of the a lecture is recorded by a computer in the lecture room. The audio is mono (single channel), and is recorded using PCM.

Consider the scenario of recording the audio and saving to the computer in the lecture room.

(a) If the sampling frequency is 44 kHz and a 16-bit sample is used, how much disk space is needed to record a 1 hour lecture? [3 marks]

Answer. *PCM will generated 16-bits, 44,000 times per second. Within 1 hour there are 3,600 seconds. Therefore a 1 hour lecture generates* $16 \times 44,000 \times 3,600$ *bits of data. That is 2,534.5 Mb or 316.8 MB per hour.*

(b) Explain both an advantage and a disadvantage of changing the sample size to 8-bits. [2 marks]

Answer. *Reducing the sample size will reduce the file size (advantage), but reduce the quality of the recorded audio (disadvantage).*

Now consider an alternative scenario, where instead of saving the the audio on a computer in the lecture room, that computer streams the audio to a central server. Each lecture room does the same: streams the live audio to a single central server. The central server, which has effectively unlimited disk space and disk write speed, saves the audio.

The computers in each lecture room and the central server are all connected via a single Ethernet switch. All links support full-duplex, 100 Mb/s.

The protocol used for streaming the audio sends UDP datagrams. Each UDP datagram contains 880 Bytes of audio data (the format of a UDP datagram is shown in the Reference Material). The UDP datagram is put into an IP datagram, which in turn is put into an Ethernet frame to be transmitted. (Only consider the Ethernet frame format in the Reference Material; ignore any overheads of the Ethernet PHY, e.g. preamble).

(c) How large is each Ethernet frame sent by a lecture room computer? [3 marks]

Answer. *With 880 Bytes of audio data, there is an 8 Byte UDP header, 20 Byte IP header and 18 Byte Ethernet header/trailer. Therefore each Ethernet frame is 926 Bytes.*

(d) With a PCM sampling frequency of 44 kHz and a 16-bit sample, how many Ethernet frames must be sent per second by a lecture room computer to achieve live streaming? [3 marks]

Answer. *Each second there is 88,000 Bytes of data generated. Each frame contains 880 Bytes of data. Therefore 100 frames per second must be sent.*

(e) How many lectures can use this system at the same time? That is, how many lecture room computers can stream to the central server? [3 marks]

Answer. *Each lecture sends 100 frames per second, and 926 Bytes per frame. That is a sending rate of 740,800 bits per second. The Ethernet switch to central server supports a data rate of 100 Mb/s. Therefore 134 lectures can be streamed at the same time.*

Question 5 [17 marks]

Consider A sending data to B across a link using flow control. A always has data available to send to B. Each DATA frame takes t_d μs to transmit and contains 80% original data and 20% header. Each ACK frame takes t_a μs to transmit. Propagation delay of the link is p μs from A to B and p μs from B to A. Assume no processing delays.

(a) With the values of $t_d = 100$, $t_a = 2$ and $p = 200$, what is the maximum efficiency possible if using stop-and-wait flow control? [3 marks]

Answer. *With stop-and-wait flow control, each DATA frame takes a total time of total μs to transfer:*

$$total = t_d + p + t_a + p$$

Within that time, 80% of t_d is spent sending original data. Therefore the efficiency is:

$$\eta = \frac{80\% \times t_d}{total}$$

With the values given, that is 15.9%.

(b) Assuming you cannot change the link data rate or propagation delay, explain two ways in which the efficiency of the data transfer can be increased while still using stop-and-wait flow control. [2 marks]

Answer. *Reduce the header, relative to the original data; reduce the ACK size; Increase the original data per DATA frame*

(c) If the ratio of original data to total DATA frame size is d, write an equation to calculate efficiency of stop-and-wait flow control, η, for any value of d, t_d, t_a and p. [3 marks]

Answer.

$$\eta = \frac{d \times t_d}{t_d + p + t_a + p}$$

(d) With the values of $t_d = 100$, $t_a = 2$ and $p = 200$, what is the maximum efficiency possible if using sliding-window flow control with a window size of 3? [3 marks]

Answer. *The total time is still 502 μs. However 3 DATA frames are delivered in this time, taking 300 μs for transmission, but 240 μs for transmission of original data. Therefore the efficiency is 240/502 = 47.8%.*

(e) What is the optimal window size when using the values of $t_d = 100$, $t_a = 2$ and $p = 200$? Explain your answer (e.g. why a smaller and larger window would be sub-optimal). The window can be any integer (it doesn't have to be related to a power of 2). [3 marks]

Answer. *A window is 6 allows A to be transmitted the window when it receives the first ACK back (at time 502 μs). A smaller window means A will spend some time waiting, therefore will be inefficient. A larger window means B requires more buffer space, but the efficiency will not be higher.*

(f) If the ratio of original data to total DATA frame size is d, write equations to calculate efficiency of sliding-window flow control, η, for any value of d, t_d, t_a, p and the window W. [3 marks]

Answer. *If* $t_d + p + t_a + p > W \times t_d$*:*

$$\eta = \frac{W \times d \times t_d}{t_d + p + t_a + p}$$

Else:

$$\eta = d$$

Question 6 [10 marks]

Consider datagram packet switching being used to send n packets across a single path with h hops, where $n >> h$. Each packet has a length of l bits. Each link has a data rate of b bits per second. Each link has a propagation delay of p seconds. Although there is processing delay and queuing delay at each node, they are both so small that you can assume they are both 0.

(a) With the values of $h = 4$, $n = 20$, $l = 1,000$, $b = 1,000,000$, and $p = 0.0001$, calculate the total time, $T_{datagram}$, it takes from when the source host initiates the data transfer until the destination host has received all of the data. [3 marks]

Answer. *See answer to next part for general equation.* $T_{datagram} = 0.0234$ *seconds.*

(b) Write an equation that gives the total time, $T_{datagram}$, it takes from when the source host initiates the data transfer until the destination host has received all of the data. [3 marks]

Answer. $T_{datagram} = n \times \frac{l}{b} + h \times p + (h-1) \times \frac{l}{b}$

(c) If all the conditions above are the same except virtual circuit packet switching was used instead of datagram packet switching, then would the total time, $T_{virtualcircuit}$, be greater than $T_{datagram}$, less than $T_{datagram}$, equal to $T_{datagram}$ or there is insufficient information to know. Explain your answer. [2 marks]

Answer. $T_{virtualcircuit}$ *would be greater than* $T_{datagram}$ *because in virtual circuit packet switching the data is delivered in the same manner as datagram packet switching, but there is an additional time before the data transfer for the virtual circuit to be established.*

(d) Explain an advantage of virtual circuit packet switching compared to datagram packet switching. [2 marks]

Answer. *Virtual circuit packet switching ensures all packets in a data transfer will follow the same path. One advantage of this is that the packets will be delivered in order, hence no need for re-ordering at destination. Another advantage is that resources can be allocated along the path to improve the performance for the data transfer.*

Question 7 [12 marks]

Consider A sending data to B across a link using stop-and-wait ARQ. A always has data available to send to B. Each DATA frame takes t_d μs to transmit and contains 80% original data and 20% header. Each ACK frame takes t_a μs to transmit. Propagation delay of the link is p μs from A to B and p μs from B to A. A uses a timeout values of TO μs, where the timer starts immediately *after* the transmission of an entire data frame (i.e. after the last bit is transmitted). Assume no processing delays.

Assume the values of $t_d = 100$, $t_a = 2$, $p = 200$, and $TO = 800$.

(a) How long does it take from the start of the transmission of a new DATA frame until the end of the reception of an ACK if the new DATA frame is lost (but the retransmission is not lost)? [3 marks]

Answer. *The total time is:*

$$T = t_d + to + t_d + p + t_a + p$$

That is, 1,402 μs*.*

(b) If A can change the timeout value, then what is the smallest value of TO that it should use? [2 marks]

Answer. *The timeout should be large enough for the ACK to be returned, i.e. 402* μs*.*

(c) Assume you determine that 90% of new packets transmitted will be delivered successfully, 9% of new packets transmitted will be lost and require a single retransmission, and 1% of new packets transmitted will be lost and the first retransmission will also be lost (but the second retransmission will not be lost). What is the average efficiency of stop-and-wait ARQ if there is a very large number of data frames to send? Use the original values, that is: $t_d = 100$, $t_a = 2$, $p = 200$, and $TO = 800$. [4 marks]

Answer. *If there are no errors the time to transfer 1 DATA frame successfully is:*

$$t_d + p + t_a + p = 502$$

If there is 1 error, the time is:

$$t_d + TO + t_d + p + t_a + p = 1402$$

If there are 2 errors, the time is:

$$t_d + TO + t_d + TO + t_d + p + t_a + p = 2302$$

With weightings of 90%, 9% and 1% respectively, the average time is 601 μs*. Therefore the average efficiency is* $0.8 \times 100/601 = 13.3\%$.

(d) Go-Back-N and Selective-Reject are two other ARQ protocols, which are based on sliding-window. Explain how Go-Back-N and Selective-Reject operate differently when after the source sends a window of frames, one of those frames is lost. [2 marks]

Answer. *Go-Back-N requires the source to retransmit all frames sent since the lost frame (as well as the lost frame). Selective-Reject requires the source to only retransmit the lost frame.*

(e) What is an advantage of Go-Back-N (when compared to Selective-Reject)? [1 mark]

Answer. *Simple at the receiver; no need to buffer frames*

Question 8 [6 marks]

(a) Draw the TCP/IP protocol stack, giving the names of each layer, for a computer (PC or laptop) attached to a wired LAN. Also, next to each layer, write the primary protocol used in that layer when web browsing on the computer (you may either give the full protocol name or the acronym). [3 marks]

Answer. *Application - HTTP*

Transport - TCP

Network - IP

Data Link - Ethernet, IEEE 802.3

Physical - Ethernet, IEEE 802.3

(b) Draw the TCP/IP protocol stack, giving the names of each layer, for a router that is attached to a wired LAN on one interface and a wireless LAN on the other interface. Also, next to each layer, write the primary protocol used in that layer when the router is the default router of the computer which is web browsing in part (a) (you may either give the full protocol name or the acronym). [3 marks]

Answer. *Network - IP*

Data Link 1 - Ethernet ; Data Link 2 - WiFi, IEEE 802.11

Physical 1 - Ethernet; Data Link 2 - Wifi, IEEE 802.11

Reference Material

Selected well-known ports:

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

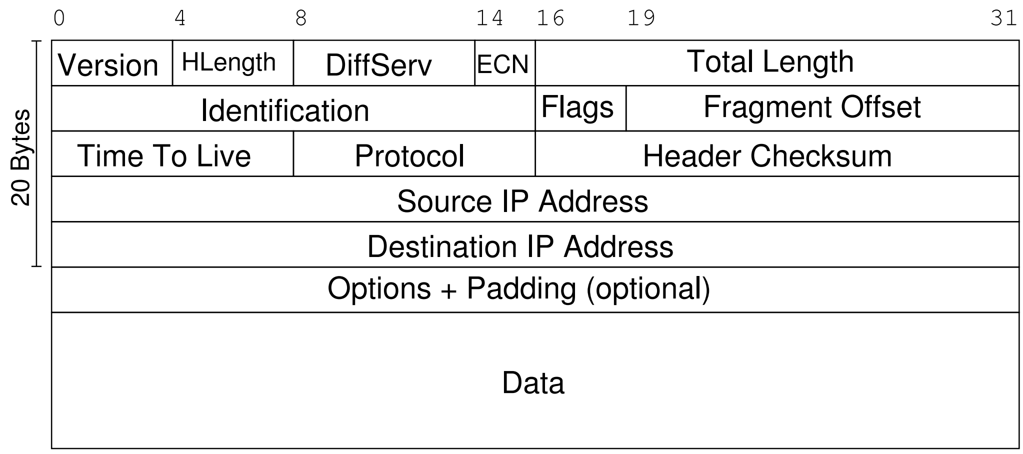

Figure 3: IP Datagram Format. Flags: Reserved, Don't Fragment, More Fragments

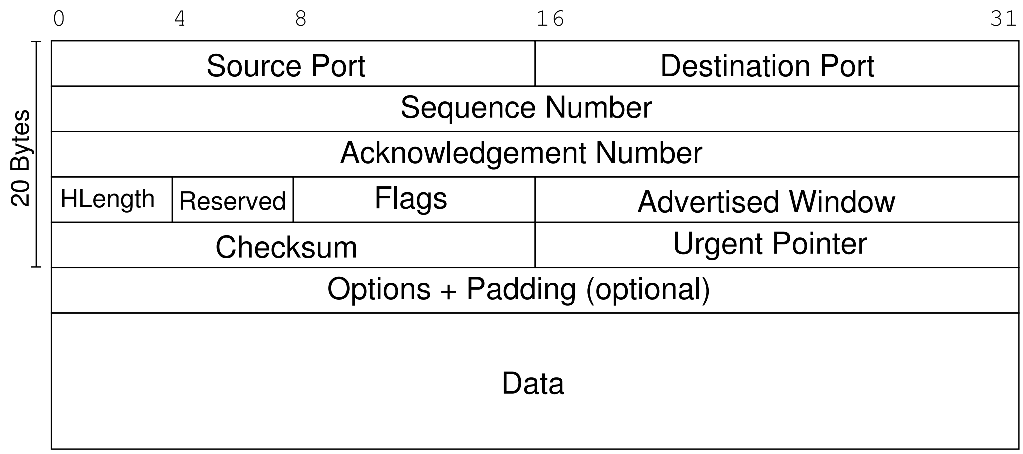

Figure 4: TCP Segment Format. Flags: CWR, ECE, URG, ACK, PSH, RST, SYN, FIN

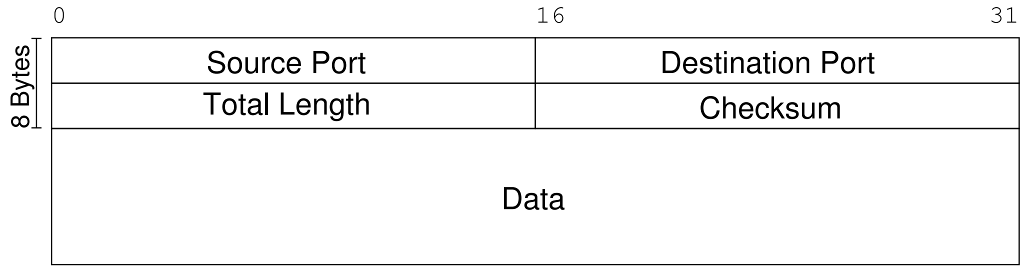

Figure 5: UDP Datagram Format

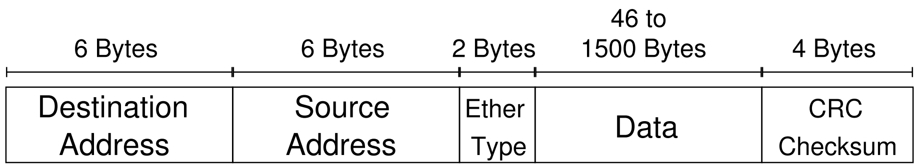

Figure 6: Ethernet Frame Format

Selected Protocol numbers:

- 1 ICMP
- 6 TCP
- 17 UDP