EECS 489 HW-1

Problem 1:

1. eecs.umich.edu looks up a.math.mit.edu
    1. eecs.umich.edu server queries the root NS for the .edu NS
    2. The root NS replies with the IP of the .edu NS
    3. The eecs.umich.edu server queries the .edu NS for IP address of mit.edu
    4. The .edu NS replies with the IP of mit.edu NS (MIT NS)
    5. The eecs.umich.edu server queries the MIT NS for the IP of math.mit.edu
    6. The MIT (NS) replies with the IP address of math.mit.edu
    7. The eecs.umich.edu server queries math.mit.edu for the IP of a.math.mit.edu
    8. The math.mit.edu NS replies with the IP of a.math.mit.edu

2. eecs.umich.edu looks up b.eecs.mit.edu
    1. eecs.umich.edu queries mit.edu (MIT NS) for the IP of eecs.mit.edu
    2. mit.edu (MIT NS) replies with address of eecs.mit.edu
    3. The eecs.umich.edu server queries the eecs.mit.edu NS for the address of b.eecs.mit.edu
    4. eecs.mit.edu replies with the IP of b.eecs.mit.edu

3. eecs.umich.edu looks up c.math.mit.edu
    1. The eecs.umich.edu server queries the math.mit.edu NS for the address of c.math.mit.edu
    2. The math.mit.edu NS replies with the IP of c.math.mit.edu

4. d.math.mit.edu looks up c.math.mit.edu
    1. d.math.mit.edu queries the local DNS server for the IP of c.math.mit.edu
    2. The local DNS server queries root NS for the IP of the .edu NS
    3. The root NS replies with the IP of the .edu NS
    4. The local DNS server queries the .edu NS for the IP of mit.edu NS (MIT)
    5. The .edu NS replies with the IP of the mit.edu (MIT) NS

6. The local DNS server queries mit.edu NS for the IP of math.mit.edu
7. The mit.edu NS replies with the IP of math.mit.edu
8. The local DNS server queries math.mit.edu for the IP of c.math.mit.edu
9. The math.mit.edu replies with the IP of c.math.mit.edu
10. The local DNS server returns the IP of c.math.mit.edu

5. d.math.mit.edu looks up b.eecs.mit.edu
    1. d.math.mit.edu queries the local DNS server for the IP of c.math.mit.edu
    2. The local DNS server queries the mit.edu for the IP of eecs.mit.edu
    3. The mit.edu NS replies with the IP of eecs.mit.edu
    4. The local DNS server queries the eecs.mit.edu NS for the IP of b.eecs.mit.edu
    5. The eecs.mit.edu NS replies with the IP of b.eecs.mit.edu
    6. The local DNS server replies with the IP of b.eecs.mit.edu

The MX type records can be used for having the same domain name for the mail and the web server. The MX type record holds as the NAME field, the domain name of the web server and the VALUE field is the canonical domain name of the mail-server.

Problem 2

From host A to router X, the packets are of size 1500 bytes (with 20 bytes of header and 1480 bytes of data). At router X, the packets are fragmented. Each 1500 byte packet is broken down into 2 packets. The first is of length 996 bytes (with 20 bytes header data and 976 bytes of data). This is so that the chinks of data are divisible by 8 (see pg 340 of Kurose). The other packet is of length 524 bytes (with 20 bytes header and 504 bytes of data).

For packets from A to X there is no fragmentation so the following are the relevant packet fields:

16 bit identifier: i (where i is the packet number)
Flag: 0 (no fragmentation)
Offset: 0

For packets from X to Y, there is fragmentation, so each packet is broken down into two packets (A and B). The relevant fields for the two packets are shown below:

Packet A:

16 bit identifier: i (where i is the identifier of the original packet)
Flag: 1 (signifies fragmentation)
Offset: 0

Packet B:

16 bit identifier: i (where i is the identifier of the original packet)
Flag: 0 (last fragment)
Offset: 122 (as a multiple of 8)

The total number of packets received at B are $2*(4 * 10^6)/(1480) = 5406$

Problem 3:

Advantages of packet switching:

1. Highly scalable
2. Less strain on hardware resources (a circuit is not reserved for every connection)
3. Cheap to implement overall
4. End point intelligence
5. Etc.

Advantages of circuit switching:

1. Better quality in communication (because of reserved resources)
2. Reliable (no losses)
3. etc.

Problem 4:

1. Three different ways:
   a. The amount of data transmitted could be protocol dependent. The advantage of this approach is that it is very simple. The disadvantage is that the protocol will be very hard to extend and scale.
   b. The end of the data could be marked with a delimiter. The advantage is that this design is extensible and scalable. The disadvantage is that the delimiting character cannot be a part of the message.
   c. The length of the message could be signified by a fixed size number (eg 4 byte big endian integer) at the beginning of the message. The advantage is that this design is also very flexible and scalable. The disadvantage is that the size of the message is limited by the size of the integer.
2. listen() only needs be called with SOCK_STREAM and SOCK_SEQPACKET. Not all protocols need to call listen().
3. accept() only needs to be called when using TCP. UDP communication does not require that accept() be called. Also accept is a blocking system call.

4. Alternate answers to 2 and 3 could be related to flexibility in program design and more fine grained control over the behavior of the program.

Problem 5:

Subnet 1: 223.1.1.0/25 (The last 7 bits cover 128 hosts)
Subnet 2: 223.1.1.128/26 (The last 6 bits cover 64 hosts)
Subnet 3: 223.1.1.192/26 (The last 6 bits cover 64 hosts)

Problem 6 (see next page)

z		h	d			h	d			h	d			h	d
	u	-	∞		u	v	10		u	v	10		u	v	10
	v	v	5		v	v	5		v	v	5		v	v	5
	x	x	2		x	x	2		x	x	2		x	x	2
	y	y	10		y	x	3		y	x	3		y	x	3
	z	z	0		z	z	0		z	z	0		z	z	0

x		h	d			h	d			h	d			h	d
	u	u	13		u	u	13		u	z	12		u	z	12
	v	-	∞		v	z	7		v	z	7		v	z	7
	x	x	0		x	x	0		x	x	0		x	x	0
	y	y	1		y	y	1		y	y	1		y	y	1
	z	z	2		z	z	2		z	z	2		z	z	2

y		h	d			h	d			h	d			h	d
	u	-	∞		u	v	12		u	v	12		u	v	12
	v	v	7		v	v	7		v	v	7		v	v	7
	x	x	1		x	x	1		x	x	1		x	x	1
	y	y	0		y	y	0		y	y	0		y	y	0
	z	z	10		z	x	3		z	x	3		z	x	3

v		h	d			h	d			h	d			h	d
	u	u	5		u	u	5		u	u	5		u	u	5
	v	v	0		v	v	0		v	v	0		v	v	0
	x	-	∞		x	z	7		x	z	7		x	z	7
	y	y	7		y	y	7		y	y	7		y	y	7
	z	z	5		z	z	5		z	z	5		z	z	5

v		h	d			h	d			h	d			h	d
	u	u	0		u	u	0		u	u	0		u	u	0
	v	v	5		v	v	5		v	v	5		v	v	5
	x	x	13		x	x	13		x	v	12		x	v	12
	y	-	∞		y	v	12		y	v	12		y	v	12
	z	z	∞		z	v	10		z	v	10		z	v	10

The maximum number of iterations is C-1 (C if you account for the last iteration to confirm a stable routing table).

This can be proved either inductively or intuitively. In a synchronous version of this protocol, the case where it takes the maximum number of iterations to stabilize the routing table is when the router at one end of the longest path holds some route updates that provide better routes to the router at the second end. These updates will take C-1 iterations to propagate, when C is the length of the longest path.

Problem 7

IP addresses of the three hosts behind NAT:
10.0.0.1
10.0.0.2
10.0.0.3

Internal IP address of the Gateway: 10.0.0.4

NATing table for the TCP connections:

Internal IP:Port	External IP:Port
10.0.0.1:10001	128.15.88.65:10024
10.0.0.1:10002	128.15.88.65:10025
10.0.0.2:10003	128.15.88.65:10026
10.0.0.2:10004	128.15.88.65:10027
10.0.0.3:10005	128.15.88.65:10028
10.0.0.3.10006	128.15.88.65:10029