ELEX 4550 : Wide Area Networks
2014 Winter Session

Solutions to Assignment 3

Question 1

(a) each SONET frame contains whole multiples of bytes; it uses byte-oriented framing.

(b) The frame rate is 8 kHz to match the sampling rate of PCM digitized speech.

(c) In each 810-byte SONET frame two bytes (A1 and A2) or 16 bits are used only for framing.

The B1 bytes detect errors for each physical link(section), the B2 bytes are used for error detection in a specific STS-1 signal and the B3 byte allows error detection for a specific payload. Error detection can then be used to compute error rates on for each of these elements.

Since the frame rate is 8 kHz, the data rate consumed by framing overhead is $2 \times 8 \times 8$ kHz $=$128 kb/s and the data rate consumed by *each* of the error detection bytes is $1 \times 8 \times 4$ kHz $=$64 kb/s.

(d) $3 \times 9 = 27$ bytes of the 810 SONET frame bytes are overheard. The overhead is thus $\frac{27}{810} = 1/30 = 3.3\%$ of the whole frame.

An STS-3c frame consists of 3 STS-1 payloads concatenated into the payload portion of a STS-3 frame. The 27 bytes of line overhead for the second and third are not required. Thus the overhead could be considered to be $\frac{1}{3}$ as much or about 1%. However, to keep things simple, the hardware transmit these bytes filled with with dummy data and the data rate at the physical layer remains unchanged. Thus the overhead at the physical layer is unchanged (3%). Either answer can be considered correct.

Question 2

The H1,H2 pointer field is an offset into the 783 byte SPE portion of the frame and is incremented or decremented to compensate for differences between the SONET and payload clock rates. Since the pointer can only take on 783 different values (0 to 782) it will "wrap around" after being incremented or decremented 783 times. If this happens every 24 hours, then an extra 783×8 bits were received in $24 \times 60 \times 60$ seconds. The bit clock frequency difference causing the shift in pointer value is thus $\frac{783 \times 8}{24 \times 60 \times 60} \approx 72$mHz ($10^{-3}$ Hz).

Question 3

Following the ALL segmentation process, we need to split the 100-byte packet into 48-byte ATM cells. There will be $\lfloor 100/48 \rfloor = 2$ full packets holding $2 \times 48 = 96$ bytes and the remaining $100 - 96 = 4$ bytes will go into the last cell. The first 5 bytes of each cell are the cell headers which will each contain a 4-bit GFC field (0x0), 8-bit VPI (50 decimal = 0x32), 16-bit VCI (82 decimal = 0x0052), 3-bit type and CLP fields (0x0) and a checksum assumed to be 0xff. Each header except the last will thus contain: 0x03, 0x20 0x08, 0x20, 0xff. The third bit of the type field will be 1 in the last frame so the fourth byte will be 0x22.

The last cell will have to be padded out to a length of 40 bytes, leaving 8 byte for the AAL5 trailer. So there will be 4 byte of data plus $40 - 4 = 36$ bytes of padding (assumed 0xaa).

The trailer will have a 16-bit reserved field (set to zero, 0x00 0x00), a 16-bit length field with value 100 (0x00 0x64) and a 4-byte CRC (4 bytes, each assumed to be 0xff 0xff 0xff 0xff).

Putting this all together we get:

```
Cell 1:

0x03, 0x20 0x05, 0x20, 0xff

48 X 0x00


Cell 2:

0x03, 0x20 0x05, 0x20, 0xff
```

```
48 X 0x00


Cell 3:

0x03, 0x20 0x05, 0x22, 0xff

0x00 0x00 0x00 0x00

36 X 0xaa

0x00 0x00 0x00 0x64 0xff 0xff 0xff 0xff
```

Question 4

A call rate λ of 5 calls/minute (only half of the calls require external trunks) multiplied by a mean call duration H of 4 minutes gives a traffic intensity of $a = 20$ Erlangs (calls). Looking at the Erlang-B curves on page 4 of the Lecture 14 notes we see that at the intersection of the horizontal line for a blocking probability of 2% and the vertical line for an offered load of 20 falls between the curves for 25 and 30 trunks so it would be reasonable to provision the switch with 27 or 28 trunks.

Question 5

PPP encapsulation requires adding: flags (0x7e) before and after the frame; a fixed header consisting of 0xff, 0x03, protocol byte of (given as 45 = 0x2d in the question); the data and FCS (assumed 0x00 0x00). Also, each flag or escape character needs to be escaped by sending an escape character with the next character xor-ed with 0x20. The data given in the question contains both flag and escape characters that would be converted into the pairs 0x7d 0x5e and 0x7d 0x5d respectively. So the values transmitted would be:

```
0x7e  0xff 0x03 0x2d

0x20 0x7d 0x5d 0x00 0x7d 0x5e 0x7d 0x5e 0xff

0x00 0x00  0x7e
```

Question 6

RFC 1234 can be found at:

https://tools.ietf.org/html/rfc1234

and describes it's purpose as: "*This memo describes a method of encapsulating IPX datagrams within UDP packets so that IPX traffic can travel across an IP internet.*"

Question 7

The IP packets consists of a header with 5 32-bit words (20 bytes) followed by the 76 bytes of data. The IPv4 packet has the following fields in its header:

- version and header length: 0x45
- optional type of service: 0x00
- total length: 96 = 0x00 0x60
- fragmentation-related fields: 0x00 0x00 0x00 0x00
- TTL: 64 = 0x40
- protocol: 17 = 0x11
- header checksum: 0x?? 0x??
- source address: 192.168.0.10 = 0xc0 0xa8 0x00 0x0a
- destination address: 192.168.0.1 = 0xc0 0xa8 0x00 0x01
- options: not present

The header contents are thus:

```
0x45 0x00 0x00 0x60
0x00 0x00 0x00 0x00
0x40 0x11 0x?? 0x??
0xc0 0xa8 0x00 0x0a
0xc0 0xa8 0x00 0x01
```

which is followed by the 76 bytes of data.

Question 8

(a) The netmask for a /17 network in binary is 17 ones followed by $32 - 17 = 15$ zeros. In hex this is: 0xff 0xff 0x80 0x00 which is 255.255.128.0.

(b) The netmask 255.255.224.0 in hex is 0xff 0xff 0xe0 0x00 which contains 19 leading zeros. So the netmask is /19 and the network address is 123.45.96.0/19.

Question 9

And-ing in the most significant 15 bits of the IP address 140.15.205.114 gives 140.14.0.0 which is in the 140.14.0.0/15 network. So yes, the address is 'in' the specified network.

Question 10

(a)

(i) Routing Tables for H.

destination	mask	gateway	interface	purpose
10.0.0.0	255.0.0.0	*	10.0.0.2	route traffic for 10.0.0.0/8 to G
0.0.0.0	0.0.0.0	192.168.12.254	192.168.0.2	gateway other traffic to R
(ii) Routing Tables for R.				
10.0.0.0	255.0.0.0	*	10.0.0.1	route traffic for 10.0.0.0/8 to G
192.168.0.2	255.255.255.255	*	192.168.12.254	traffic for host 192.168.0.2 to H
0.0.0.0	0.0.0.0	10.6.1.2	10.0.0.1	gateway other traffic to G
(iii) Routing Tables for G.				
10.0.0.1	255.255.255.255	*	10.6.1.2	route all traffic for 10.0.0.1 to R
10.0.0.2	255.255.255.255	*	10.0.0.3	route all traffic for 10.0.0.2 to H

(b) G has no route for destination address 142.232.15.23. If a router has no route to a destination it will have to ignore (or "drop") it.

To prevent this from happening again, the destination should inform the source that it received a frame that it could not forward. This is done by sending an ICMP type 3 ("destination unreachable") packet.

The arrival of the un-routable packet could have been cause by transient reasons (e.g. temporary link failure or router reconfiguration) so it may not immediately cause the source to stop sending packets to that destination. These ICMP messages are mainly used by network managers to detect and resolve configuration problems.