1. Suppose a TCP message that contains 2048 bytes of data and 20 bytes of TCP header is passed to IP for delivery across two networks of the internet (i.e., from the source host to a router to the destination host). The first network uses 14-byte headers and has an MTU of 1024 bytes and the second uses 8-byte headers with an MTU of 512 bytes. Each network's MTU gives the size of the largest IP datagram that can be carried in a link-layer frame. Give the sizes and offsets of the sequence of fragments delivered to the network layer at the destination host. Assume all IP headers are 20 bytes **(10 Marks)**

Ans:

. Consider the first network. Packets have room for $1024 - 14 - 20 = 990$ bytes of IP-level data; because 990 is not a multiple of 8 each fragment can contain at most $8 \times \lfloor 990/8 \rfloor = 984$ bytes. We need to transfer $2048 + 20 = 2068$ bytes of such data. This would be fragmented into fragments of size 984, 984, and 100.

Over the second network (which by the way has an illegally small MTU for IP), the 100-byte packet would be unfragmented but the 984-data-byte packet would be fragmented as follows. The network+IP headers total 28 bytes, leaving $512 - 28 = 484$ bytes for IP-level data. Again rounding down to the nearest multiple of 8, each fragment could contain 480 bytes of IP-level data. 984 bytes of such data would become fragments with data sizes 480, 480, and 24.

2. Suppose an IP packet is fragmented into 10 fragments, each with a 1% (independent) probability of loss. To a reasonable approximation, this means there is a 10% chance of losing the whole packet due to loss of a fragment. What is the probability of net loss of the whole packet if the packet is transmitted twice. **(10 Marks)**

(a) Assuming all fragments received must have been part of the same transmission?

(b) Assuming any given fragment may have been part of either transmission?

(c) Explain how use of the Ident field might be applicable here.

Ans:

(a) The probability of losing both transmissions of the packet would be $0.1 \times 0.1 = 0.01$.

(b) The probability of loss is now the probability that for some pair of identical fragments, both are lost. For any particular fragment the probability of losing both instances is $0.01 \times 0.01 = 10^{-4}$, and the probability that this happens at least once for the 10 different fragments is thus about 10 times this, or 0.001.

(c) An implementation *might* (though generally most do not) use the same value for Ident when a packet had to be transmitted. If the retransmission timeout was less than the reassembly timeout, this might mean that case (b) applied and that a received packet

might contain fragments from each transmission.

3. (a) IP currently uses 32-bit addresses. If we could redesign IP to use the 6-byte MAC address instead of the 32-bit address, would we able to eliminate the need for ARP? Explain why or why not? **(5 Marks)**

(b) Suppose hosts A and B have been assigned the same IP address on the same Ethernet, on which ARP is used. B starts up after A. What will happen to A's existing connections? Explain how "self ARP" (querying the network on startup for one own IP address) might help with this problem.

(5 Marks)

Ans:

2. The answer is no in practice, but yes in theory. MAC address is statically assigned to each hardware. ARP mapping enables indirection from IP addresses to the hardware MAC addresses. This allows IP addresses to be dynamicaly reallocated when the hardware moves to the different network. So using MAC addresses as IP addresses would mean that we would have to use static IP addresses.

Since the Internet routing takes advantage of address space hierarchy (use higher bits for network addresses and lower bits for host addresses), if we would have to use static IP addresses, the routing would be much less efficient. Therefore this design is practically not feasible.

After B broadcasts any ARP query, all stations that had been sending to A's physical address will switch to sending to B's. A will see a sudden halt to all arriving traffic. (To guard against this, A might monitor for ARP broadcasts purportedly coming from itself; A might even immediately follow such broadcasts with its own ARP broadcast in order to return its traffic to itself. It is not clear, however, how often this is done.)

If B uses self-ARP on startup, it will receive a reply indicating that its IP address is already in use, which is a clear indication that B should not continue on the network until the issue is resolved.

4. Suppose a router has built up the routing table shown below. The router can deliver packets directly over interfaces 0 and 1, or it can forward packets to routers R2, R3 or R4. Describe what the router does with a packet addressed to each of the following destinations:

(a) 128.96.39.10

(b) 128.96.40.12

(c) 128.96.40.151

(d) 192.4.153.17

(e) 192.4.153.90

SubnetNumber	SubnetMask	NextHop
128.96.39.0	255.255.255.128	Interface 0
128.96.39.128	255.255.255.128	Interface 1
128.96.40.0	255.255.255.128	R2
192.4.153.0	255.255.255.192	R3
⟨Default⟩		R4

Table 4.14 Routing table for Exercise 21.

Ans:

Apply each subnet mask and if the corresponding subnet number matches the SubnetNumber column, then use the entry in Next-Hop. (In these tables there is always a unique match.)

(a) Applying the subnet mask 255.255.255.128, we get 128.96.39.0. Use interface0 as the next hop.

(b) Applying subnet mask 255.255.255.128, we get 128.96.40.0. Use R2 as the next hop.

(c) All subnet masks give 128.96.40.128 as the subnet number. Since there is no match, use the default entry. Next hop is R4.

(d) Next hop is R3.

(e) None of the subnet number entries match, hence use default router R4.

5. Consider the network in the figure below, using link-state routing. Suppose the B-F link fails, and the following then occur in sequence:

(a) Node H is added to the right side with a connection to G

(b) Node D is added to the left side with a connection to C

(c) A new link D-A is added

The failed B-F link is now restored. Describe what link-state packets will flood back and forth. Assume that the initial sequence number at all nodes is 1, and that no packets time out, and that both ends of a link use the same sequence number in the LSP for that link, greater than any sequence number either used before.

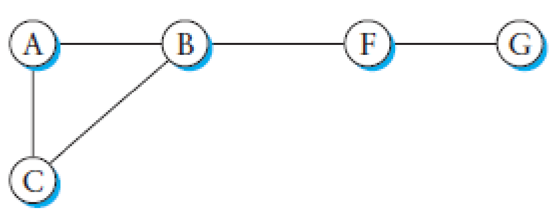

Ans:

'. There is some confusion in the last paragraph of this exercise. OSPF routers send out *one* LSP, with one sequence number, that describes all the router's connections; however, the language "both ends of a link use the same sequence number in their LSP for that link" incorrectly suggests that routers send out a different LSP (or at least different LSA) for each link, each with its own sequence number. While it is certainly possible for link-state routing to take this approach, it is not how OSPF works and it is not what the text describes. We will use OSPF-style numbering here. We will also assume that each node increments its sequence number only when there is some change in the state of its local links, not for timer expirations ("no packets time out").

The central *point* of this exercise was intended to be an illustration of the "bringing-up-adjacencies" process: in restoring the connection between the left- and righthand networks, it is not sufficient simply to flood the information about the restored link. The two halves have evolved separately, and full information must be exchanged.

Given that each node increments its sequence number whenever it detects a change in its links to its neighbors, at the instant before the B—F link is restored the LSP data for each node is as follows:

node	seq#	connects to
A	2	B,C,D
B	2	A,C
C	2	A,B,D
D	2	A,C
F	2	G
G	2	F,H
H	1	G

When the B–F link is restored, OSPF has B and F exchange their full databases of all the LSPs they have seen with each other. Each then floods the other side's LSPs throughout its side of the now-rejoined network. These LSPs are as in the rows of the table above, except that B and F now each have sequence numbers of 3. (Had we assigned separate sequence numbers to each individual link, every sequence number would be 1 except for link B—F.)

The initial sequence number of an OSPF node is actually $-2^{31}+1$.