Networks Sample Exam Solutions

(1-a) Three important architectural principles for computer network software are *hierarchical modularity*, *encapsulation*, and *distributed scripts*. Briefly explain these principles in the context of layered computer network architectures.

Hierarchical modularity: network software is structured as a hierarchy of layers

- each layer offers certain *services* to the higher layers, while hiding from the higher layers the *details* of how those services are implemented

Encapsulation:

- at the sender, layer N may add control information to the data it receives from layer N+1 before passing the (increased) data to layer N-1; at the receiver, layer N-1 passes data to layer N, which can read, act upon, & remove this control information before passing the (reduced) data up to layer N+1
- each layer should not need to know which portion of the upper layer's data is control information, or its meaning

Distributed Scripts: *peer entities* (at same layer, but in 2 different computers) are programmed as if data transmission were "horizontal", even though actual communications are "vertical" (except in the physical medium)

- together, these peer entities execute *distributed scripts*

Networks Sample Exam Solutions

(1-b) 200 nodes are connected to a 1,500 metre length of coaxial cable. Using some protocol, each node can transmit 50 frames/second, where the average frame length is 2,000 bits. The transmission rate at each node is 100 Mbps (where 1 Mbps = 1,000,000 bps). What is the *efficiency* of this protocol?

Per-node throughput = 50 frames/second = (50×2,000) bits/second
= 100,000 bits/second

System throughput = (200×50) frames/second = (200×100,000) bits/second
= 10,000 frames/second = 20,000,000 bits/second
= 20 Mbps

Maximum rate = (100,000,000/2,000) frames/second = 100 Mbps
= 50,000 frames/second

Efficiency = (10,000/50,000) = (20/100) or (20,000,000/100,000,000)
= 0.2, or 20%

Networks Sample Exam Solutions

(1-c) The Internet was originally intended for robust transfer of computer-to-computer data over long distances. Briefly explain why connectionless packet-switching was preferred to circuit-switching in the IP layer.

- Combination of reasons, including:
  - no set-up delay
  - no blocking
    - of course, there is no guarantee that any data reaches its destination
  - flexibility in transmission bit-rates – in contrast with circuit-switching, which is usually tied to a few pre-determined bit-rates
  - no "path" $\Rightarrow$ more reliable *(route around problems)*
  - more efficient use of network resources when traffic is bursty
    - in circuit-switching with bursty traffic:
      - could set up a new circuit for each burst
      - could hold original circuit for duration of data transfer
      - BUT: both of these solutions are wasteful of network resources

Networks Sample Exam Solutions

(1-d) The Hamming distance between 2 Datalink layer codewords is defined to be the number of bit positions in which the codewords differ. Briefly explain how this is used in the General Parity Check error-handling scheme, mentioning the limitations of the scheme for error detection and error correction.

- for any error detection/correction scheme, we can define the minimum Hamming distance (or "minimum distance") of the scheme as the ***smallest number*** of bit errors that changes one valid codeword into another
  - if the way of computing the check bits is known, a list of all the valid codewords can be compiled and stored at the receiver
  - when a word W is received, the receiver finds the ***closest valid codeword to W (in Hamming distance)*** and takes this codeword as the transmitted codeword

Limitations: if the minimum distance of an error-handling scheme is D, this scheme can **detect** any combination of ≤ D-1 bit errors and **correct** any combination of strictly less than D/2 bit errors

(alternatively – if you want to detect B bit errors, use a scheme with minimum distance at least B+1; if you want to correct B bit errors, use a scheme with minimum distance at least 2B+1)

Networks Sample Exam Solutions

(2-a) Consider a Data Link Layer with the following measured parameters:

- **frame transmission time at the sender is TRANSF = 200 microseconds**
- **ACK or NAK transmission time at the receiver is TRANSA = 4 microseconds**
- **link propagation delay is PROP = 20 microseconds**
- **frame processing time at sender and receiver is 0 (in other words, negligible)**
- **overall round-trip probability of frame error on the link is r = 0.02**

Assume that for both the Stop-and-wait and Go-back-n ARQ schemes, the TIMEOUT at the sender is chosen optimally. The average packet throughput in each scheme is given by the following formulas:

$$\text{throughput}_{SW} = (1-r)/(\text{TRANSF}+\text{TIMEOUT})$$

$$\text{throughput}_{GBN} = (1-r)/(\text{TRANSF}+(r \times \text{TIMEOUT}))$$

If you want to ensure an average packet throughput of at least 4,500 packets/second, which of these ARQ schemes could you use? Justify your answers mathematically.

Networks Sample Exam Solutions

TIMEOUT is chosen optimally, so TIMEOUT = TRANSA + 2×(PROP+PROC)

= 4 + 2×(20+0)

= 44 microseconds

Since TRANSF = 200 microseconds and r = 0.02, we can now calculate

$throughput_{SW} = (1 - 0.02) / (200\times10^{-6} + 44\times10^{-6})$ = **4,016.39 packets/second**

$throughput_{GBN} = (1 - 0.02) / (200\times10^{-6} + (0.02\times44\times10^{-6}))$ = **4,878.53 packets/second**

Since we want to ensure a throughput of at least 4,500 packets/second, Go-back-n can be used under these circumstances, but not Stop-and-wait.

Networks Sample Exam Solutions

(2-b) Draw timing diagrams to show how a Go-back-n ARQ scheme copes with

1. a damaged data frame;

2. a lost data frame; and

3. a lost ACK.

The following examples were used in lectures, and something like these timing diagrams is perfectly acceptable to answer this question.

Go-back-n ARQ, damaged data frame

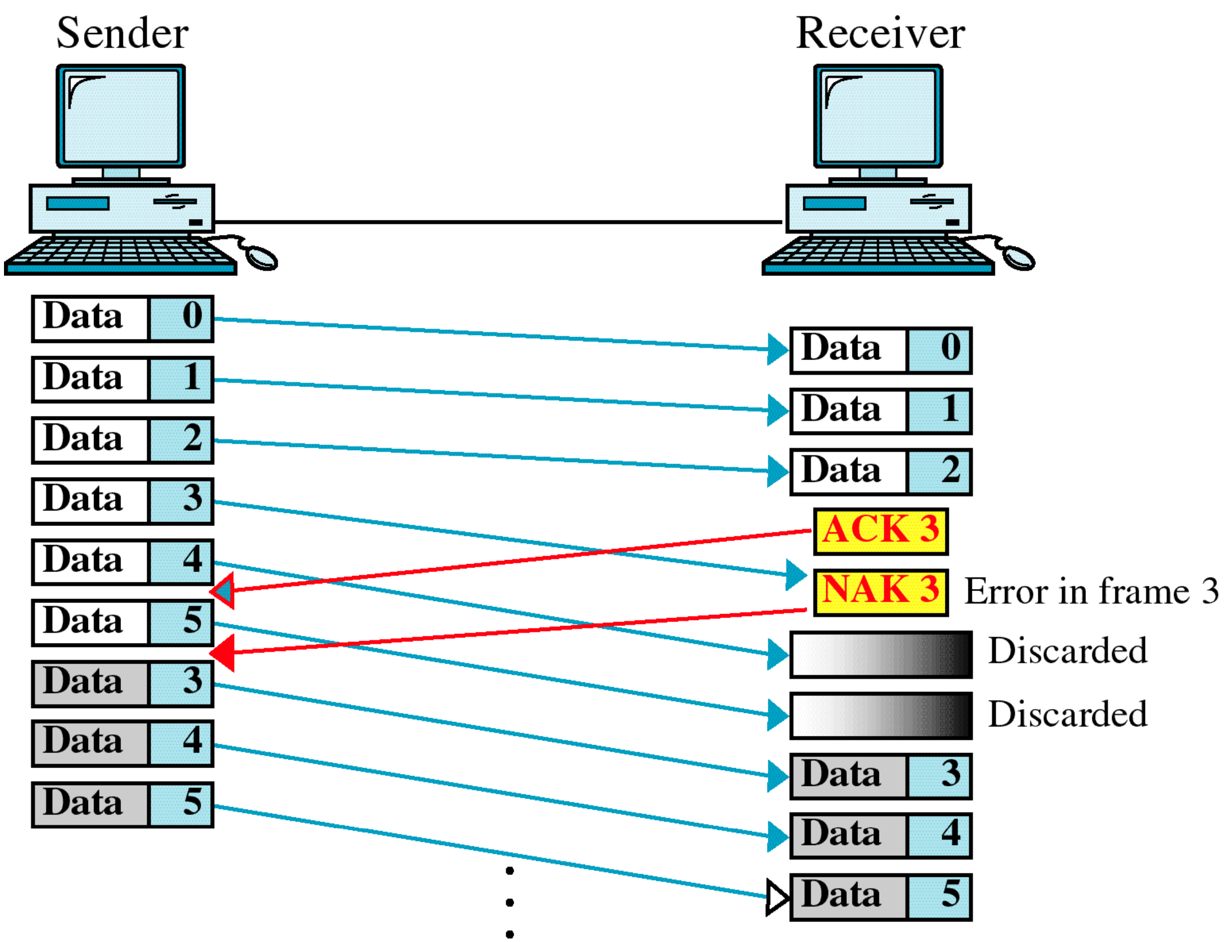

Go-back-n ARQ, lost data frame

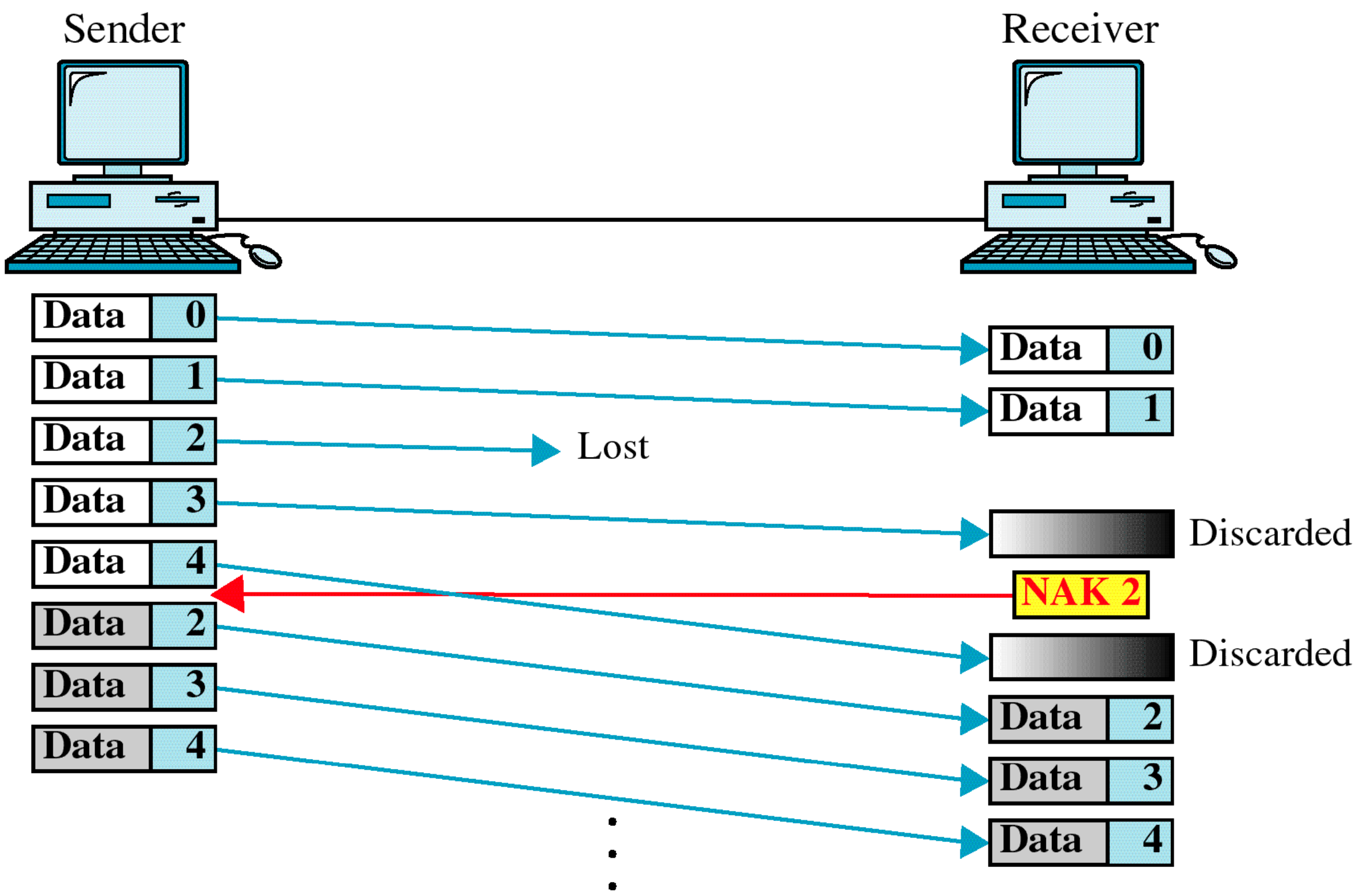

Go-back-n ARQ, lost ACK

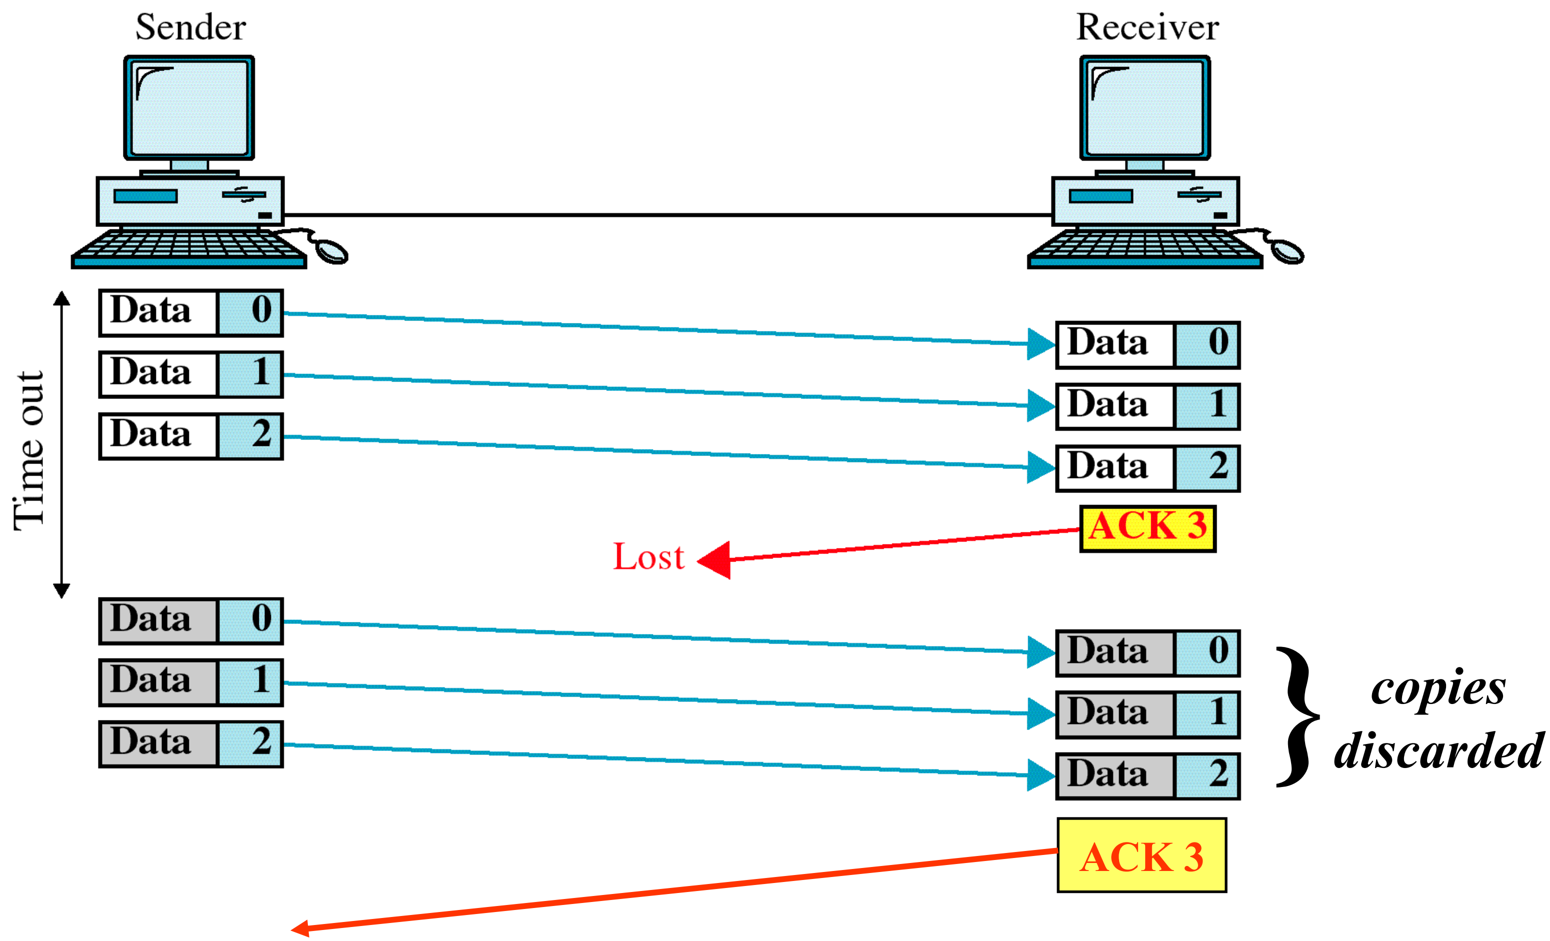

Networks Sample Exam Solutions

(3-a) The throughput of an IEEE 802.5 Token Ring can be determined by the formula

```
throughput = 1/(TRANSF + (TRANSF/THT)×PROP)
```

where `PROP` is the one-way channel propagation delay, `TRANSF` is the average frame transmission time, and `THT` is a constant value of 10 milliseconds. Using this formula, state and explain the effect on this Token Ring's throughput of the following changes:

the length of the channel is <u>increased</u> (everything else held constant);

Channel length increased $\Rightarrow$ PROP increased $\Rightarrow$ denominator increased $\Rightarrow$ *throughput decreased*

the average frame length is <u>increased</u> (everything else held constant).

Frame length increased $\Rightarrow$ TRANSF increased $\Rightarrow$ denominator increased $\Rightarrow$ *throughput decreased*

Networks Sample Exam Solutions

(3-b) Briefly describe *circuit switching* and *virtual circuit packet switching*, mentioning their principal differences.

Circuit-switching: a path is set up in the network between the sender and the receiver, by making the appropriate connections in network switches

- the necessary network resources are reserved for the connection prior to any data transfer; if this is not possible, the connection request is blocked
- these reserved resources are then held for the duration of the connection, regardless of actual usage

Virtual circuit packet switching: a route is set up in the network between sender and receiver, by making appropriate entries in network routers' routing tables

- resources may or may not be reserved for this route. If resources need to be reserved and are not available, the connection request is blocked

Principal differences:

- circuit switching creates a path in the network; virtual circuit p.s. creates a route which exists only in software
- In circuit switching, the links in the path cannot be shared during the connection; in virtual circuit p.s. they can
- Also, c-s is ideal for "smooth" network traffic, whereas v-c is ideal for bursty traffic (especially if it's long-lived &/or interactive)

Networks Sample Exam Solutions

(4-a) Consider a TCP connection using the slow-start congestion control scheme with an initial THRESHOLD value of 64 kB and a Maximum Segment Size (MSS) of 4 kB. The receiver's advertised window is initially 24 kB. The first transmission attempt is numbered 0, and all transmission attempts are successful except for Timeouts on attempt number 4. In the ACKs for transmission attempt number 9 and subsequently, the receiver's advertised window is reset to 20 kB.

Find the size in kB of the *sender's congestion window* for its first 11 transmission attempts (numbers 0 – 10).

Transmission no.	Sender's Congestion Window (kB)	Threshold (kB)	Rcvr. Window (kB)
0	4	64	24
1	8	64	24
2	16	64	24
3	24	64	24
4	24	64	24
5	4	12	24
6	8	12	24
7	12	12	24
8	16	12	24
9	20	12	24
10	20	12	20

Networks Sample Exam Solutions

(4-b) In IP-based networks, a sending host can find the physical address which corresponds to the IP address of its intended destination by using the Address Resolution Protocol (ARP). Briefly explain how ARP works.

- each host maintains an *ARP cache* of mappings between IP addresses and physical addresses
  - ARP cache entries *time out* after (typically) 15 minutes
- if destination's IP address not in sending host's ARP cache, sending host broadcasts *ARP query* on local network
  - ARP query asks for physical address for destination's IP address
  - ARP query includes sending host's IP address and physical address (so other hosts can enter this in their ARP caches)
- if destination is a host on local network, it sends its physical address to sending host in an *ARP reply*
- if destination not on the local network, ARP server replies with its own physical address: *proxy ARP*

Networks Sample Exam Solutions

(5-a) In this diagram, A, B, and C are routers. The ovals represent LANs, labeled with their network ID… The routers are using DISTANCE-VECTOR routing.

1. Show the *initial routing tables* exchanged by the routers.

initial routing tables:

A:

Network ID	*Cost*	*Next Hop*
1	*1*	--
2	*1*	--

B:

Network ID	*Cost*	*Next Hop*
1	*1*	--
3	*1*	--

C:

Network ID	*Cost*	*Next Hop*
2	*1*	--
3	*1*	--

Networks Sample Exam Solutions

2. Show how router A *updates* its routing table if it first receives B's initial routing table.

When A receives B's initial routing table, since B is 1 hop from A, A modifies B's table as follows:

Modified B:

Network ID	Cost	Next Hop
1	2	B
3	2	B

The only entry in common is for Network 1, and A's original entry is lower-cost, so A's new table is

A:

Network ID	*Cost*	*Next Hop*
1	*1*	--
2	*1*	--
3	*2*	*B*

Networks Sample Exam Solutions

(5-b) Suppose instead that LINK-STATE routing is being used. The following link costs have been determined…

1. Show the *link-state packets* each router floods to all other routers.

A:

Advertiser	Network	Cost	Neighbour
A	1	1	B
A	2	4	C

B:

Advertiser	Network	Cost	Neighbour
B	1	3	A
B	3	1	C

C:

Advertiser	Network	Cost	Neighbour
C	2	1	A
C	3	2	B

Networks Sample Exam Solutions

2. Show all the steps used by router A to determine its *shortest-path spanning tree* after it has received link-state packets from all other routers.

The sequence of nodes being made Permanent is: A, Net 1, B, Net 3, C, Net 2.

A's shortest-path spanning tree and link-state routing table:

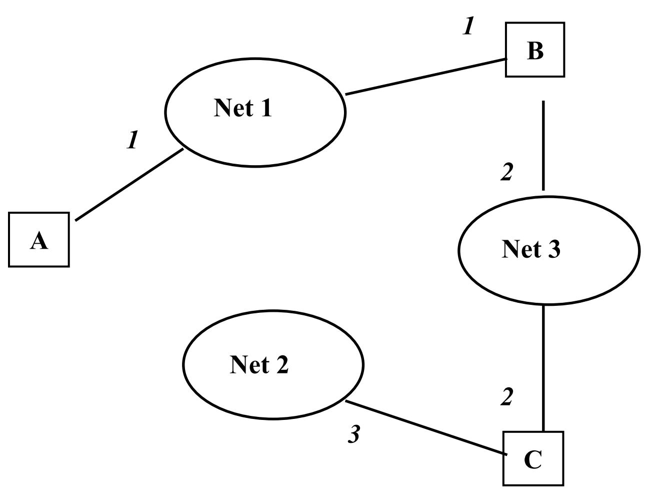

Net	Cost	Next Router
1	1	--
2	3	B
3	2	B