



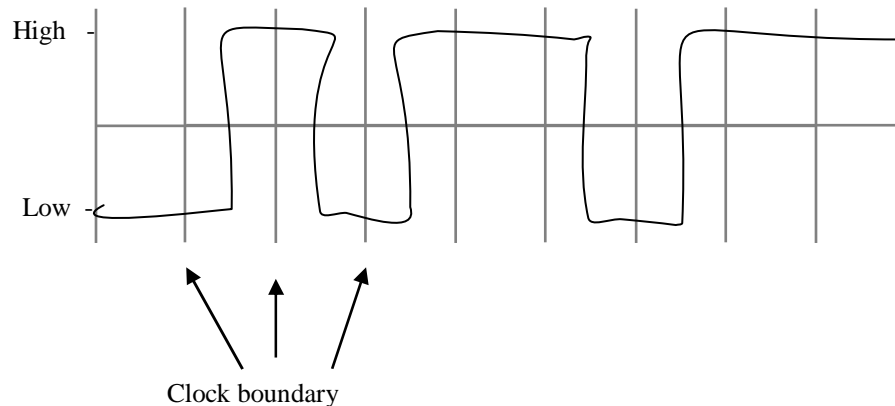
## 1. Encoding, Framing and Error Control

A host wishes to send the following nine bits:

**011101100**

Assume that NRZI signals start low.

(a) Show the NRZI signal encoding of this data. [1 point]



(b) True or false (identify all true statements): [2 points for circling all correct answers, -0.5 points for each wrong answer] Correct answers in *italics*

- CRC can both correct and detect errors.
  - No, it can only correct.
- *2-d parity can detect all 2-bit errors and correct all 1-bit errors.*
- 4B/5B encoding solves the problem of long sequences of zeroes but still has issues with long sequences of 1s.
  - No. 4B/5B uses NRZI to get over the issue of a long series of 1s.
- *When bandwidth is plenty and errors are rare, error checking may actually add unnecessary overhead.*


If you interpret the “overhead” as bandwidth overhead, then it can be argued that error checking will not add any overhead. The intended interpretation was “computational overhead” though.

(c) Identify one key advantage that SONET framing offers over bit/byte stuffing protocols. [1 point]


*Fixed vs. variable length frames. SONETs fixed frames provide several advantages, including simplicity, high speed, ability to multiplex higher or lower speed links etc.*

## 2. Bridging

Consider hosts X, Y, Z, W, U and learning bridges B1, B2, B3, B4 with initially empty forwarding tables, connected in the following manner:



(a) Assuming B1 has the lowest bridge identifier, and B4 the highest, show the spanning tree for the above network. [2 points]



(b) Say B4 is connected to U, B2, B3 and Z via its ports 1, 2, 3 and 4, respectively. Assuming no failures, which of these ports will B4 *never* forward on (i.e. which ports are *blocked*)? [1 point]

3, 4

(c) Suppose X sends its first packet to Z. Which bridges learn where X is? Does Y's network interface see this packet? [1 point]

*All bridges will learn. Yes, Y's interface will see this packet.*

(d) Suppose Z now sends its first packet to X. Which bridges learn where Z is? Does Y's network interface see this packet? [1 point]

***B1, B2, B3 will learn. Y will not see this packet.***

(e) Suppose Y now sends to X. Which bridges learn where Y is? Does Z's network interface see this packet? [1 point]

***B1, B2 will learn. Z will not see this packet.***

### 3. Packets and Addressing

(a) Host A is on LAN 1 and host B is on LAN 2. The two LANs are interconnected by a router. The MTU of the two LANs are 1000B and 500B respectively. Suppose an application on host A executes 5000 writes to an application on host B. Each write results in one MTU sized IPv4 packet on LAN 1. Each packet consists of an IP header (20B) and the data. The IP headers do not carry any options. How many packets from host A to host B traverse LAN 2? What are the sizes of the packets? [2 points]

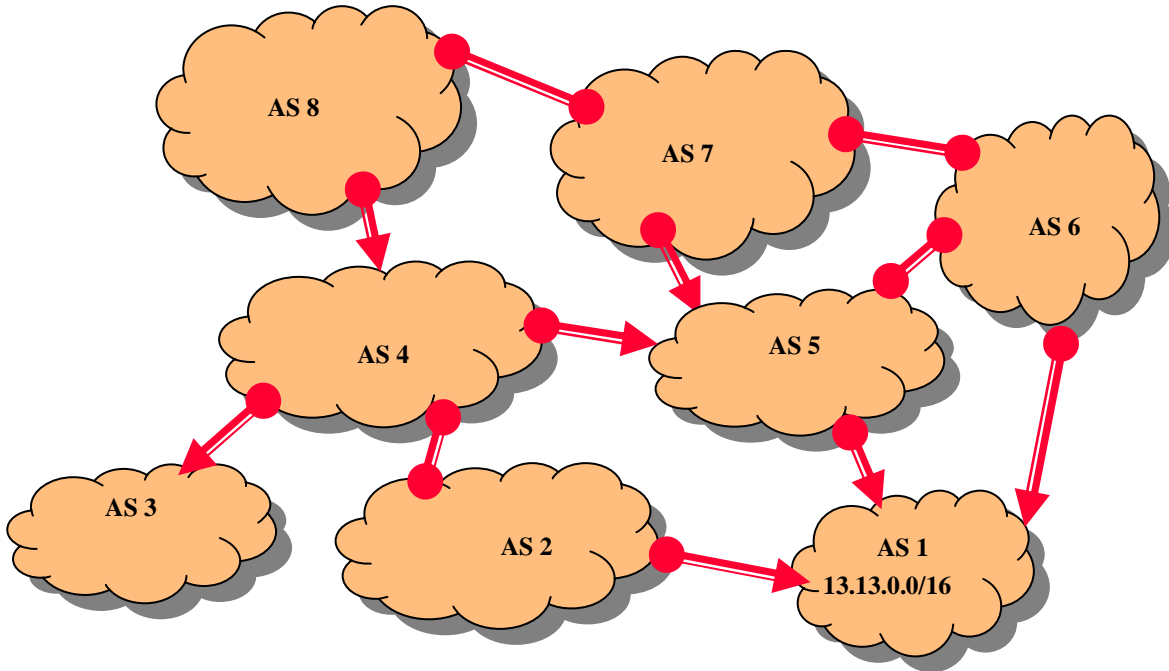

Each IP packet has 20B of header and 980B of payload.

*The first fragment will be 500B long, with 20B of header and 480B of payload. The second fragment will have the same header/payload sizes. This leaves 20B of payload left, which will go into the final fragment, making the fragment size 40B. In all, there will be 15,000 packets on LAN 2.*

(b) Which of the following prefixes *are contained in* the CIDR prefix 201.10.0.0/21 (circle all)? [2 point]

- **201.10.4.0/24**
- 201.10.8.0/23
- 201.10.24.0/22
- **201.10.6.0/23**

## 4. BGP and Policy



Consider the AS-level topology in the picture above. Say AS 1 starts by announcing its prefix, 13.13.0.0/16, into the network at time 0. At time 1, each of AS 1's neighbors – AS 2, AS 5 and AS 6 – propagate the route to their neighbors. And so on.

All routing messages exchanged in this network pertain to the prefix 13.13.0.0/16.

- (a) Show the path vector that AS 5 sends out at *time 1*. Which of its neighbors does AS 5 send this path vector to? [2 points]

Say,  $P = 13.13.0.0/16$

*Path vector:*  $\{P, (5, 1)\}$

*Sent to:* 4, 6, 7

(b) Show the path vector that AS 7 sends out at *time 2*. Which of its neighbors does AS 7 send the path vector to? [2 points]

*{P, (7, 5, 1)}* (note that this is preferred over (7, 6, 1) due to prefer-customer routing)

*Sent to: 6, 8*

(c) A *valid* AS level path is one that is allowed by economic policies. List all possible, valid, AS level paths from AS 8 to the prefix 13.13.0.0/16. Is *AS8--AS7--AS6--AS1* a valid AS level path from AS 8 to the prefix 13.13.0.0/16. Why or why not? [3 points]

*Valid paths: 8—4—5—1 and 8—7—5—1*

*8—7—6—1 is not a valid path since it violates “valley-freeness” (two consecutive peering links).*

(d) Which path will AS 8 use to route to the prefix 13.13.0.0/16? Why? [2 points]

*It will use 8—4—5—1 due to prefer customer routing.*

## 5. Routing and Routers

(a) Which of the following is true about unicast routing (circle all correct answers): [2 points for circling all correct answers, -0.5 points for each wrong answer.]

- In an RIP network which has no cycles with more than 3 edges in it, split-horizon and poisoned reverse are sufficient to prevent the count-to-infinity problem.
  - Count-to-infinity is still possible here. See class notes.
- The convergence time of RIP depends only on the number of routers and the number of links in the network.
  - Also depends on link costs.
- The router-level Internet path from A to B will always be the exact inverse of the router-level path from B to A.
  - No. Due to BGP policies, MED, hot-potato routing, etc.
- *The BGP route between two hosts may not be the shortest possible in terms of the number of router hops.*

(b) For each of the following statements, identify whether they are true or false. Provide a one line explanation for your answer. Full credit only for the correct answer *and* the correct explanation. [3 points]

- All router activities related to TTL decrement are always handled in the fast path of the router (i.e., in special purpose silicon chips on the input ports) and the slow path of the router (i.e. the router processor) is never involved.

*False. When TTL goes to zero, CPU must generate ICMP error.*

- Copies of router forwarding table are kept at all input ports.

*True. This enables fast distributed look-ups and prevents route processor from becoming a central bottleneck.*

- IP address lookup in a router takes fewer clock cycles than Ethernet MAC address lookup in a switch (assume that the lookup tables have the same total number of entries).


*True. For the same length tables, IP address lookup will take lesser time since IP addresses are hierarchical. This is despite the fact that IP lookup needs longest prefix match, masking etc. Also, the fact that Ethernet addresses are long and require an exact match makes Ethernet address lookup cumbersome.*



(c) A certain router B in a network receives the following link state updates from other nodes in the same network:

| Update from | Sequence# | Neighbor/cost  |
|-------------|-----------|----------------|
| C           | 5         | B/8, A/4       |
| A           | 20        | B/2, C/4, G/10 |
| C           | 10        | B/10, A/4      |
| F           | 50        | B/1, G/6       |
| G           | 100       | F/6, A/10      |

(c.1) Show the network topology, with nodes, edges and link costs, as constructed by B after it receives the above link state packets. [2 points]



(c.2) Suppose that the A—B link goes down. Say, A and B quickly recompute shortest paths to other network nodes, but their link state updates are yet to reach their neighbors. Is there a possibility of packets looping in the network? Why? [2 points]

*Yes. There is a possibility of looping. After A and B do the recomputation, the shortest A—B path is through C, but C still thinks that the shortest C—B path is through A. So packet will be transiently stuck in an A—C—A... loop.*

(d) Outline one advantage and one disadvantage of MOSPF compared to DVMRP. [2 points]

*Advantage: No state required where there are no receivers.*

*Disadvantage: May have to keep per-receiver state even when there are no senders.*

(e) From a router's point of view, name one challenge in supporting IP multicast efficiently. [1 point]

*Look-up on source+group address becomes cumbersome. Also, multicast routing and forwarding are a bit more challenging to run than their unicast counterparts.*

## 6. Internet Architecture

(a) Outline two disadvantages and one advantage of circuit switching compared to packet switching. [3 points]

*Disadvantages: When resources are scarce, idle circuits could waste resources. Second, circuit setup and teardown is complex.*

*Advantage: Easy to provide QoS and isolation.*

(b) Identify all the correct answers: [2 points for circling all correct answers, -0.5 points for each wrong answer]

- In the TCP/IP architecture, layering must be strictly followed (e.g. applications cannot run directly over IP).
  - No, tcpdump (or ethereal) is an example application that violates layering.
- *TCP's reliability mechanisms make it a poor choice for time-sensitive applications such as voice.*
  - *True. Wait for retransmissions could impose unreasonable delays.*
- All applications that require reliable data delivery must run on top of TCP.
  - No. Can implement reliability at the app-layer, even on top of UDP.
- *The use of NATs goes against the design principle of "fate sharing".*
  - True. When a NAT box fails, all connection state is lost and end-hosts must restart from scratch.