# Review Materials: Set A

## EE122 Fall 2012

2011 WS1 + WS2 + HW3Q5, now with biodegradable packaging

## 1.1 Acronyms

Consider the following acronyms: A, AIMD, BGP, CDN, CIDR, CNAME, DHCP, DNS, ECN, FIN, HTTP, ICANN, LPM, MSS, MTU, MX, NAT, NS, PMTU, PTR, RFC, RIP, RST, RTO, SSTHRESH, STFU, SYN, TCP, TLD, UDP.

Match the acronyms to the descriptions below, using each acronym exactly once. Supply exactly one answer for each item below. In some cases, an acronym might fit more than one question, but there is only one way to answer these so that every acronym is used exactly once. For terms that aren't familiar to you, use the book or the web.

i) A mechanism that allows the network to indicate congestion to TCP without having to drop the packet. **_ECN**

ii) The DNS record type that provides the IP address of a given hostname. **A**

iii) The DNS record type that provides the canonical name associated with a given hostname. **CNAME**

iv) The DNS record type that provides the authoritative name server associated with a given domain. **NS**

v) The DNS record type that provides the mail server associated with a given domain. **MX**

vi) The DNS record type that provides a hostname associated with a given IP address **PTR**

vii) The name of the message used to initiate a TCP connection. **SYN**

viii) A message (control flag) used to terminate a TCP connection abruptly. **RST**

ix) A message (control flag) used to terminate a TCP connection smoothly. **FIN**

x) The kind of window adjustment algorithm used in TCP. **AIMD**

xi) The size of the largest IP packet that can be sent end-to-end along a path without fragmentation. **PMTU**

xii) The largest sized segment that can be sent by TCP. **MSS**

xiii) The largest size IP packet that can be sent across a particular link. **MTU**

xiv) The most common examples of this acronym are .com, .net, and .org. **TLD**

xv) The documents used by the IETF to describe protocol standards. **RFC**

xvi) An in-line network device that rewrites IP addresses (and often transport ports). **NAT**

xvii) A widely used routing protocol that does not necessarily compute lowest-cost paths. **BGP**

xviii) The organization that is ultimately responsible for address allocation and domain name management. **ICANN**

xix) A method for matching packets to routing entries that determines which among multiple matching entries determines the packet's destination. **LPM**

xx) The length of time TCP waits for an acknowledgement before timing out and initiating a retransmission. **RTO**

xxi) The most widely used reliable transport protocol. **TCP**

xxii) An unreliable transport protocol. **UDP**

xxiii) The protocol that provides newly arrived hosts with their own IP addresses (and also gives them other information). **DHCP**

xxiv) An IP addressing scheme that does not involve the traditional A, B, and C address classes. **CIDR**

xxv) Akamai is an example of one. **CDN**

xxvi) The TCP congestion control state variable that determines when to switch from slow-start to congestion avoidance. **SSTHRESH**

xxvii) An application-level protocol that typically runs over TCP. **HTTP**

xxviii) An application-level protocol that typically runs over UDP. **DNS**

xxix) A distance-vector routing protocol. **RIP**

xxx) The appropriate response to Stanford students when they bring up the 2010, 2011 or 2012 Big Games. **STFU**

## 1.2 A TCP Sob Story

Due to budget cuts, the CS department is forced to buy cheaper networking equipment. It turns out that one of the switches has a defective port that has a strange dropping behavior. On a particular port, it drops every fourth packet, in both the sending and receiving directions. To be precise, if one only looks at the packets being received, it drops the $4^{th}$, $8^{th}$, $12^{th}$, ... packet. Similarly, if one only looks at the packets being sent, it drops the $4^{th}$, $8^{th}$, $12^{th}$, etc.

This problematic port is the port that attaches the department chair's computer, host A, to the network, so every fourth packet sent to A is dropped, and every fourth packet sent by A is dropped.

Suspecting a problem, the CS department runs a test by initiating a file transfer from host A to an offsite host B. The RTT of the packets is 1 msec, the packet transmission time is .01msec, and the RTO (timeout value) in the TCP connection is 100msec (you won't need these exact numbers, they just set the relative magnitude of the various timescales). Host A uses a transport protocol with a sliding window flow control with a constant window size of 5 packets and a duplicate ACK threshold of 3 packets (3 duplicate ACKs lead to a retransmission). The test consists of sending a series of data packets (labeled D1, D2, …) from A to B.

We want to describe the transmissions *from* A (by writing Dx for the appropriate x). The first five data packets are shown, in addition to the starting SYN and ACK packets, just to get you started (and just to clarify, we assume that the packet count, in terms of dropping, starts with the SYN; therefore, in these first few packets, D2 and only D2 is dropped). The packets from A that will be dropped are marked in bold. Please fill in the remaining 15 packets sent from A. (And remember that the packets coming *to* A are also dropped in a similar one-every-four pattern, with the packet count starting with the SYN-ACK, so the first packet dropped in the receiving direction is the ACK sent in response to D4.)

| | | | |
|---|---|---|---|
| SYN, | ACK, | D1, | **D2,** |
| D3, | D4, | D5, | **D6,** |
| D2 (TO), | D7, | D8, | **D9,** |
| D10, | D6 (TO), | D11, | **D12,** |
| D13, | D9 (TO), | D14, | **D15,** |
| D16, | D12 (TO) | | |

How many fast retransmits occur during the transmission of this set of packets? 0

How may timeouts occur during the transmission of this set of packets? 4

## 1.3 TCP Congestion Control Details

In the following, assume that the MSS=1000 bytes and that in all computations CWND is rounded down to the nearest integer.

During the congestion avoidance phase, TCP updates the window following each ACK by the following equation:

$$\text{CWND} \mathrel{+}= \text{MSS} / \text{Int}(\text{CWND} / \text{MSS})$$

where Int(x) is the integer part of x.

In the slow-start phase, TCP updates the window following each ACK by

$$\text{CWND} \mathrel{+}= \text{MSS}$$

The algorithm leaves slow-start when CWND > SSThresh (note the strict inequality, which is implementation dependent rather than a property of the TCP spec.). When undergoing a Fast Retransmit, assume the TCP uses the "advanced" version discussed in the congestion control lecture (slide 38).

In the answers on the next page, denote the packet containing bytes 1401 to 1500 as the $15^{th}$ MSS, and denote the ACK expecting the next byte of 15001 as an ACK for the $16^{th}$ MSS.

Consider a TCP connection with a CWND of 12000 bytes and then undergoes a timeout. Assume all data up to 16000 have been ACKed (i.e., the last ACK's expected next byte was 16001). What is the value of CWIN, and which packets does TCP send, at each step in the following process? The first two lines are filled out, to indicate the required form of the answers.

| *ACK received or timeout* | *cwin* | *Packets sent (MSS #'s)* |
|---|---|---|
| Timeout | 1000 | 17 |
| ACK for 17th MSS | 2000 | 18, 19 |
| ACK for 18th MSS: | **3000** | **20, 21** |
| ACK for 19th MSS: | **4000** | **22, 23** |
| ACK for 20th MSS: | **5000** | **24, 25** |
| ACK for 21st MSS: | **6000** | **26, 27** |
| ACK for 22nd packet: | **7000** | **28, 29** |
| ACK for 23rd packet: | **7142** | **30** |
| ACK for 24th packet: | **7284** | **31** |
| ACK for 25th packet: | **7426** | **32** |
| ACK for 26th packet: | **7568** | **33** |
| ACK for 27th packet: | **7710** | **34** |
| ACK for 28 th packet? | **7852** | **35** |
| ACK for 29th packet: | **7994** | **36** |
| ACK for 30th packet: | **8136** | **37, 38** |
| ACK for 31st packet: | **8261** | **39** |
| ACK for 31st packet? | **8261** | |
| ACK for 31st packet: | **8261** | |
| ACK for 31st packet: | **7130** | **32** |
| ACK for 31st packet: | **8130** | |
| ACK for 31st packet? | **9130** | **40** |
| ACK for 31st packet: | **10130** | **41** |
| ACK for 31st packet: | **11130** | **42** |
| ACK for 39th packet? | **4130** | **43** |

## 1.4 HTTP and Caching

i) Consider a case where a client A is retrieving files F and G from web site B. F and G are both 125KB (i.e., one megabit). The RTT between A and B is 10msec (note, these are round-trip-times, not one-way latencies), and the bandwidth between the sites is 10Mbps. Assume all TCP SYN/ACK packets and HTTP request packets are negligible in size. How long does it take A to retrieve both files under the following circumstances:

Sequential (one-at-a-time) requests with nonpersistent TCP connections:_ **240**_

Concurrent requests with nonpersistent TCP connections: **220**_

Sequential requests within a single persistent TCP connection: **230**_

Pipelined requests within a single persistent TCP connection: **220**_

ii) Consider the same situation as in i), but assume that rather than a dedicated link there is a large shared link with many flows traversing it, and each TCP connection gets 10Mbps (adding additional flows does not significantly change the bandwidth per TCP connection, because there are thousands of flows on the link). Now, how long does it take A to retrieve both files under the following circumstances:

Sequential (one-at-a-time) requests with nonpersistent TCP connections: **240**_

Concurrent requests with nonpersistent TCP connections:_ **120**

Sequential requests within a single persistent TCP connection: **230**

Pipelined requests within a single persistent TCP connection:_ **220**

iii) Consider the same situation as in i), except that A is only downloading file F and there is now a cache C between A and B. All requests from A to B go through cache C, and assume the bandwidth along the path from A to C is 1gbps and the RTT between A and C is negligible, while the bandwidth along the path from C to B is 10mbps with an RTT of 10msec. Note, these are round-trip-times, not one-way latencies. As above, assume that the file is 125KB (i.e., one megabit) and that all TCP SYN/ACK packets and HTTP request packets are negligible in size.
Assume the cache operates as follows: (where the origin server refers to the site named in the URL)

- If the object is not in the cache, the request is forwarded to the origin server

- If the object is in the cache, and the cache entry has not timed out (i.e., the cache TTL has not expired), the object is returned to the client
- If the object is in the cache, but the cache entry has timed out, the cache issues a conditional-GET to the origin server, asking if the object has changed since this object was cached: if the origin server responds that it hasn't, the cache returns the cached object, otherwise the origin server responds with the updated object which the cache forwards to the client.

How long does it take for A to receive the file under the following circumstances:

- The file is not in the cache.___ **121**___

- The file is in the cache and the TTL has not expired. __**1**____

- The file is in the cache, the TTL has expired, but the file has not been changed. _**21**_

- The file is in the cache, the TTL has expired, and the file has changed. _**121**_

## 1.5 BGP Policy Extravaganza

i) Length-based policies:

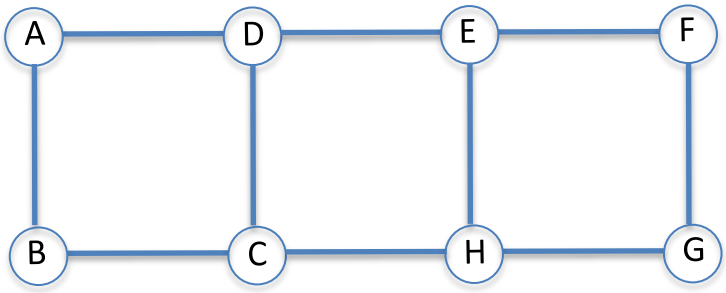


Consider the interdomain connectivity graph above, where each of the ASes A through H are running BGP. Every domain advertises its chosen routes to all its neighbors (i.e., there is no limiting export policy). Assume that in this case all the domains want shortest paths (as measured in terms of the number of ASes traversed) and, when there is a tie, they prefer to route advertised with the letter closest to the beginning of the alphabet.

What path do packets take from A to G?_**A-B-C-H-G**_
What path do packets take from G to A?_ **G-F-E-D-A**

ii) Export policies: Using the same graph, now assume that nodes only advertise routes to the vertically connected node and a horizontal peer if they have one: assume that A-D are peers, E-F are peers, and C-H are peers. When advertising routes, nodes either advertise all or none. When selecting a route, ASes prefer the shortest path and, when there is a tie, they prefer to route advertised with the letter closest to the beginning of the alphabet.

What path do packets take from A to G?_ **A-D-C-H-E-F-G**
What path do packets take from G to A?__ **G-F-E-H-C-D-A**_

iii) Nonstandard policies: Using the same graph, now assume that each domain's route preferences can be expressed in terms of their preferences over their neighbors (i.e., a domain prefers any route coming from one neighbor over any route coming from another neighbor, no matter what their relative lengths are). Assume the preferences are as follows (where x > y means the domain prefers routes from x over routes from y).

Domain A: D>B
Domain B: A>C
Domain C: H>B>D
Domain D: C>E>A
Domain E: H>F>D
Domain F: E>G
Domain G: F>H
Domain H: C>E>G

What path do packets take from A to G?_ **A-D-C-H-G**_
What path do packets take from G to A?_ **G-F-E-H-C-B-A**

Note: there is not a unique right answer to this question: BGP could reach several different fixed points. All you need to do is give an answer that is stable; that is, your answer should reflect a set of routes that, if BGP were to start with these routes as its initial condition, BGP would continue to retain that set of routes.

## 2.1 Basic Concepts and Terms

Consider the following concepts and terms: ~~Autonomous Systems~~, ~~Bellman-Ford~~, ~~Congestion Control~~, ~~Count-to-Infinity~~, ~~Cryptographic Hash~~, ~~Digital Signature~~, ~~Dijkstra's Algorithm~~, ~~Duplicate Acknowledgements~~, ~~End-to-End Principle~~, ~~Fate Sharing~~, ~~Flow Control~~, ~~Hidden Terminal~~, ~~Peer-to-Peer~~, ~~Poisoned Reverse~~, ~~Policy Oscillations~~, ~~Priority Packet Scheduling~~, ~~Slow Start~~, ~~Public Key Cryptography~~, ~~Symmetric Key Cryptography~~, ~~Three-Way Handshake~~.

Match these terms to the descriptions, using each term exactly once**.**

(a) A design style that involves many equivalent nodes, rather than a few specialized servers. **Peer-to-Peer**

(b) Something that can happen in BGP that results in unstable routing tables. **Policy Oscilations**

(c) The problem in wireless networking of a sender not being able to detect if its transmissions will collide with those of another node because the sender cannot itself hear the other node's transmissions, even though the receiver can. **Hidden Terminal**

(d) A technique in distance-vector routing that prevents some cases of looping by not allowing a node A to advertise a route to node B if A would forward packets to node B as the first hop of that route. **Poisoned Reverse**

(e) If the technique in item (d) above is not applied, this phenomenon can occur, which then requires many iterations of the routing protocol before a routing loop is removed. **Count-to-Infinity**

(f) The mechanism used by TCP to open up the congestion window quickly, and is used only if the connection is not operating in Congestion Avoidance. **Slow Start**

(g) A class of encryption algorithms that require use of a shared secret key. **Symmetric Key Cryptography**

(h) A class of encryption algorithms that allow a host to make an encryption key widely known while remaining the only entity that can decrypt messages encrypted with this key. **Public Key Cryptography**

(i) The mechanism used by TCP to reliably establish a connection. **Three-Way Handshake**

(j) The entities that BGP provides routes between; BGP describes its paths in terms of a series of these. **Autonomous Systems**

(k) A cryptographic primitive that ensures that the particular piece of data did come from

a particular source (or equivalently, prevents a source from repudiating that it sent that message). **Digital Signature**

(l) A cryptographic primitive that allows the receiver to conclude that the message (or file) was not tampered with in transit. **Cryptographic Hash**

(m) One of the basic principles of the Internet architecture, which suggests keeping functionality out of the network unless necessary. **End-to-End Principle**

(n) A basic Internet design principle that suggests that state should be stored on the nodes that care about that state, so that the only time the state is unavailable is when the node itself is not functional. **Fate Sharing**

(o) How the sender keeps from overloading the receiver in a TCP connection. **Flow Control**

(p) How the sender keeps from overloading the network in a TCP connection. **Congestion Control**

(q) Signals that allow a TCP source to assume a packet has been dropped without waiting for a timeout. **Duplicate Acknowledgements**

(r) A distributed algorithm to compute shortest paths, which is an example of a distance-vector routing algorithm. **Bellman-Ford**

(s) A centralized algorithm used to compute shortest paths, which is used in link-state routing protocols. **Dijkstra's Algorithm**

(t) A packet scheduling mechanism that allows a router to give better service to one class of packets. **Priority Packet Scheduling**

## 2.2 General Multiple Choice

i. **Five Basic Design Decisions:** As we discussed in class, the Internet architecture was shaped by five basic design decisions. Please list the two in the following list that are **NOT** among these five decisions: **b, f**

(a) Layering
(b) Longest prefix match
(c) Best-effort service
(d) The end-to-end principle and fate sharing
(e) A single universal internetworking layer
(f) Sliding window flow control
(g) Packet switching

ii. **ARP:** A typical ARP exchange goes as follows:

(a) Initiating host sends: ARP request
(b) Responding host sends: ARP response

Which of these messages are broadcast? **a**

iii. **Netmask:** Which of the following methods are ways a host can learn the netmask for the subnet? **a, d**

(a) Configuration
(b) ICMP
(c) ARP
(d) DHCP
(e) NAT

iv. **ICMP:** Which of the following are valid ICMP messages? **a, c, e, g**

(a) Need Fragmentation
(b) OMG
(c) Source Quench
(d) Invalid Address Format
(e) Host Unreachable
(g) TTL Expired

Which ICMP message (in the list above) is used in Traceroute? **g**

Which ICMP message (in the list above) is used to discover path MTU? **a**

v. **Peer-to-Peer:** P2P systems typically do some combination of three tasks, searching (e.g., keyword search), lookup (mapping name to location), and download.

Which approach is often used for download? **c**
(a) Some form of flooding
(b) Distributed Hash Tables
(c) Chunking

Which approach is typically used for search? **a**
(a) Some form of flooding
(b) Distributed Hash Tables
(c) Chunking

Which approach is typically used for lookup? **b**
(a) Some form of flooding
(b) Distributed Hash Tables
(c) Chunking

Which factor is most responsible for making chunking advantageous? **b**
(a) Number of participating peers
(b) Asymmetry of bandwidth (downloading at higher rate than uploading)
(c) Lack of centralized control
(d) Self-scaling

vi. **Cryptography:** Which one of the following is an easy way for host A to use public key cryptography to authenticate host B? **d**
(a) Encrypt B's public key with A's private key
(b) Ask B to encrypt A's private key with B's public key
(c) Ask B to encrypt a nonce using B's public key
(d) Ask B to decrypt a nonce that has been encrypted with B's public key

vii. **Security:** What are the 3 security goals (as described by the esteemed Professor Paxson)? **b**
(a) Cryptography, Isolation, Authentication
(b) Confidentiality, Integrity, Availability
(c) Cash, Infamy, Awe
(d) Consistency, Isolation, Availability
(e) Confidentiality, Integrity, Authentication
(f) Credibility, Intrusion, Anonymity

viii. **Security:** Consider the attack categories:

(a) Eavesdropping
(b) Disrupting
(c) Spoofing
(d) Scanning
(e) Injection
(f) ACK splitting
(g) Opportunistic ACKing

For each of the following attack behaviors, match them to an attack category:

Sending more ACKs than received packets: **f**
Listening to all messages sent on WiFi in a coffee shop: **a**
Sending packets with a fake source IP: **c**
Sending messages imitating someone else's ports and sequence numbers: **e**
Sending ACKs before data is received: **g**
Jamming WiFi signals in a classroom during a test (hey, that's a good idea!) **b**
Sending to an arbitrary destination and see if there is a response **d**

ix. **Software-Defined Networks:** True or False?

- SDN enables networks to do things they never could before. **F**
- Currently, networking is the least intellectually solid of the basic systems areas in CS. **T**
- None of the current vendors have endorsed SDN. **F**
- The three basic abstractions in SDN are Distributed State, Forwarding, and Specification. **T**
- The Network Operating System constructs a logical model of the network topology. **T**
- The virtualization layer presents a complicated network model to the control program. **F**
- Professor Shenker is attempting to brainwash us into liking SDN in order to make money. **F**
- Professor Shenker has two kids in college and needs the money. ~~**T**~~ **F**

## 2.3 General Short Questions

i. **TCP Throughput:** Consider two TCP connections whose throughput obeys the TCP throughput equation. The first TCP connection has the following parameters: MSS = 1000 bytes, RTT = .2msec, drop rate = .5% The second TCP connection has the following parameters: MSS = 500 bytes, RTT = .8msec, drop rate = 2%.

What is the ratio of throughputs (the throughput of the first TCP connection divided by the throughput of the second TCP connection)? **16**

ii. **Fair Shares:** Consider a shared link L with five connections. Each connection is limited, by its own access link (which it uses to reach the shared link L), to the following bandwidths:

Flow 1: 1Gbps
Flow 2: 2Gbps
Flow 3: 3Gbps
Flow 4: 4Gbps
Flow 5: 5Gbps

- If the shared link L has capacity C=10Gbps, what are the fair shares? **1, 2, 7/3, 7/3, 7/3**

- For which values of C (the capacity of the shared link) do flows 2 and 3 have the same fair shares? **$C \leq 9$** [express your answer as bounding inequalities on C]

- For which values of C (the capacity of the shared link) do flows 4 and 5 have the same fair shares? **$C \leq 14$** [express your answer as bounding inequalities on C]

- For which value(s) of C will flow 3 get 2.5Gbps? **10.5**

- If a particular flow gets less than its access bandwidth, does any flow receive more bandwidth than that flow? **No**

iii. **Transfer times:** You are trying to transfer the contents of a 1.25terabyte disk drive between here and New York, and have at your disposal two methods: (a) sending the data over a 100mbps link or (b) sending the drive by Federal Express (with a guarantee that it will arrive in 24 hours). Assume the network charges $10^{-10}$ cents per bit transmitted, and Federal Express charges $30 for the package.

- Which is faster? **Federal Express**
- Which is cheaper? **Network**

iv. **Headers:** You are accessing a web site using your browser, from a host that is connected to an Ethernet within Soda Hall. A packet sniffer on the Soda Hall Ethernet

captures a packet from your web session, which has TCP, IP, HTTP and Ethernet headers: starting from the outermost header (the header with bits at the very front of the packet), what is the order of the headers you need to traverse before reaching the payload? **Ethernet, IP, TCP, HTTP**

v. **Public Key Encryption:** You are releasing your Ph.D. thesis to the world, and want to make sure that everyone knows that it came from you, rather than from some imposter posing as you. Would you encrypt the file with your private key or with your public key? **Private**

vi. **IP Multicast:** Consider the (a) DVMRP and (b) CBT multicast routing designs. When the first packet is sent to a multicast group G, in which design is that packet broadcast? **DVMRP**

vii. **Ethernet:** Consider three hosts H1, H2, and H3 on an Ethernet; in the first seven time slots the following transmissions take place:

Slot 1: H1, H2 transmit
Slot 2: H1 transmits
Slot 3: H2, H3 transmit
Slot 4: idle
Slot 5: H1, H2, H3 transmit
Slot 6: H1, H3 transmit
Slot 7: H2 transmits

What is the backoff counter of the three hosts? (backoff counter starts at zero, is incremented to one after first backoff, etc.)

H1: **2**
H2: **0**
H3: **3**

- Assuming no new data arrives during or after slot 7 for any of the hosts, what is the probability that in slot 8 H1 will be the only host transmitting? **6/21**
- Under the same assumption, what is the probability that in slot 8 H2 will be the only host transmitting? **0**
- Under the same assumption, what is the probability that in slot 8 H3 will be the only host transmitting? **2/21**
- Assume that no host sends a packet in slot 8. What is the probability that H3 will be the only host sending in slot 9? **1/12**

viii. **Aloha:** Assume that there are N stations using the Aloha protocol, which is described by the probability p that a node with data will transmit in a given slot. If p=.5, write an expression for the average number of packets successfully sent per slot: $N2^{-N}$ (i.e., one and only one station sends)

Now let's assume that the wonderful folks in our EE department figure out how to build radios that can carry one or two packets in the channel simultaneously, but not more (i.e., if one packet is transmitted it gets through; if two packets are transmitted they both get

through; if three or more packets are transmitted, they **all** experience a collision). If p=.5, write an expression for the average number of packets successfully sent per slot: **$N^2 2^{-N}$**

ix. **Advanced Routing:**

- Is the Failure-Carrying Packets (FCP) design more appropriate for OSPF or RIP or BGP? (pick one) **OSPF**
- How does Routing-Along-DAGs (RAD) prevent loops? **By always ensuring that the routing entries form a DAG.**
- Do either FCP or RAD have the count-to-infinity problem? **No, because they automatically prevent loops.**
- Do you understand the mechanism for detecting and preventing policy oscillations? (if you answered yes, then please join my research group.) **No.**

## 2.4 Spanning Tree and Self-Learning

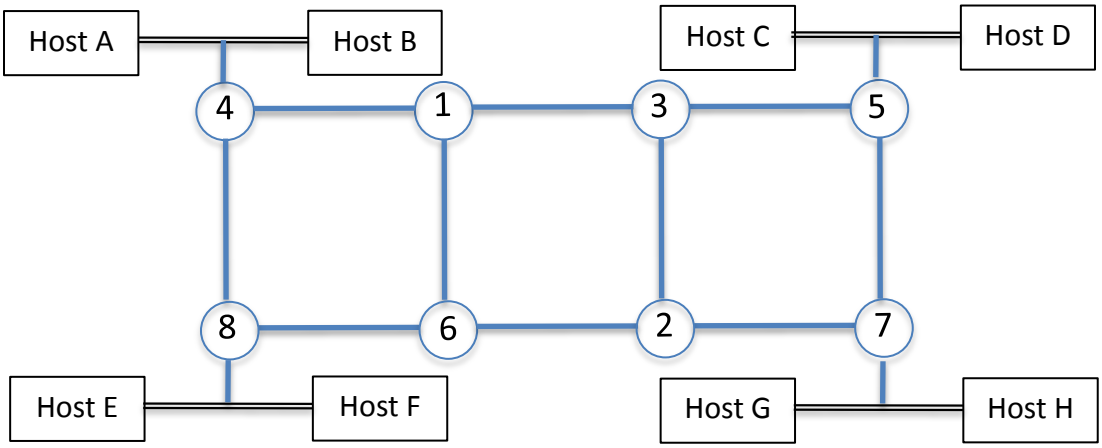


Consider the layer 2 network above connecting the switches numbered 1 through 8 (for the purposes of the spanning tree protocol, these numbers are the switch IDs). Switches 4, 8, 5, and 7 have local Ethernets that each have two hosts on them (Hosts A through H). Recall that, when constructing the spanning tree, (i) the root is the bridge with the lowest ID, and (ii) when there is more than one shortest path to the root, the path whose first hop goes through the bridge with the lower ID is chosen. Also, each link is considered to have a length of one when computing path lengths.

i. Compute the spanning tree for this network. List the links in the spanning tree (denote, for example, a link between nodes 1 and 4 as 1-4, etc.). **4-8, 4-1, 1-6, 1-3, 3-2, 3-5, 2-7**

ii. Assume that the switches are "self-learning" and consider the following set of transmissions:

(a) host A to host H
(b) host H to host E
(c) host D to host F
(d) host E to host A
(e) host C to host D
(f) host E to host H

Which of these transmissions are **NOT** broadcast to all nodes? **(d) (e) (f)**___
Which of these transmissions are dropped by their first hop switch?_ **(e)**_____

iii. In transmission (f), what path (in terms of the node IDs traversed) does the packet take from host E to host H? **8-4-1-3-2-7**_

## 2.5 Wireless

We consider three collision resolution schemes:

- Scheme X (pure carrier sense): Never send when you hear someone else transmitting, but otherwise can send whenever you want.

- Scheme Y (classic MACA): No carrier sense. Nodes wishing to communicate use an RTS-CTS-Data-ACK exchange. Nodes overhearing an RTS wait to allow the CTS to be sent. If no CTS is heard, the node can transmit. If a CTS is heard (even if no earlier RTS is heard), the node is quiet for the entire duration of the data transmission.

- Scheme Z (hybrid approach closer to 802.11): Carrier sense. Nodes overhearing *either* an RTS or a CTS are quiet for the entire duration of the transmission (data and ACK).

We have four wireless nodes A, B, C, D, where A can only hear B (but not C or D), B can only hear A and C (but not D), C can only hear B and D (but not A) D can only hear C (but not A or B). A and B are in the midst of a communication, and C has been listening to their exchange so far (and so has heard whatever RTS's or CTS's B has sent so far). While A and B are in the "sending data" part of their exchange, C decides that it wants to communicate with D. Consider two cases:

i. A is sending data to B.

- If scheme X is used, would C be allowed to send a message to D? **Yes (C would not hear A)**
- If scheme Y is used, would C be allowed to send a message to D? **No (C heard CTS)**
- If scheme Z is used, would C be allowed to send a message to D? **No (C heard CTS)**

ii. B is sending data to A.

- If scheme X is used, would C be allowed to send a message to D? **No (C hears B sending)**
- If scheme Y is used, would C be allowed to send a message to D? **Yes (C did not hear CTS)**
- If scheme Z is used, would C be allowed to send a message to D? **No (both carrier sense and hearing RTS)**

## 2.6 HTTPS

Consider a client (using a browser) interacting with a web site (based on a server) using HTTPS. Consider the following set of unordered messages, and list them in the correct order in which they would be sent, using all messages (a) - (j).

Browser to server:

(a) Set of cryptographic algorithms and protocols supported
such as (TLS+RSA+AES128+SHA1) or (SSL+RSA+3DES+MD5) or (…)
(b) TCP ACK packet
(c) User's password for site encrypted with session key K
(d) Session key K (encrypted with server's public key)
(e) TCP SYN packet

Server to browser:

(f) Agreeing to session key K
(g) TCP SYNACK
(h) Proposed set of cryptographic algorithms
such as (SSL+RSA+3DES+MD5)
(i) Server's digital certificate
(j) Response to user's password (login completed)

Message 1:_____ **e**
Message 2:_____ **g**
Message 3:_____ **b**
Message 4:_____ **a**
Message 5:_____ **h**
Message 6:_____ **i**
Message 7:_____ **d**
Message 8:_____ **f**
Message 9:_____ **c**
Message 10:_____ **j**

## 2.7 Putting it All Together

Consider the network in the diagram below, where a newly arrived laptop has just been plugged into the Ethernet. The laptop is trying to reach a file server that is located on the same Ethernet (as is the DHCP server), but the DNS server is on a remote network. Below are the initial messages **sent** and **received** by the laptop (no other messages are shown, and assume these packets are sniffed on the Ethernet to which the laptop is attached).

The MAC addresses are shortened, for convenience, and the laptop begins without an IP address, but is later assigned one (4.9.8.22) by DHCP.

The following messages are sent, in some order. Below they are listed in alphabetical order, but your job is to list them in order.

- DNS: DNS Response
- File: ARP Response
- File: TCP SYNACK to Laptop
- Laptop: ARP Request (in order to send to DNS)
- Laptop: ARP Request (in order to send to File)
- Laptop: DHCP Discovery
- Laptop: DHCP Request (aka “DHCP Accept”)
- Laptop: DNS Request
- Laptop: TCP SYN to File
- Laptop: TCP ACK to File
- DHCP: DHCP ACK
- DHCP: DHCP Offer
- Router: ARP Reply

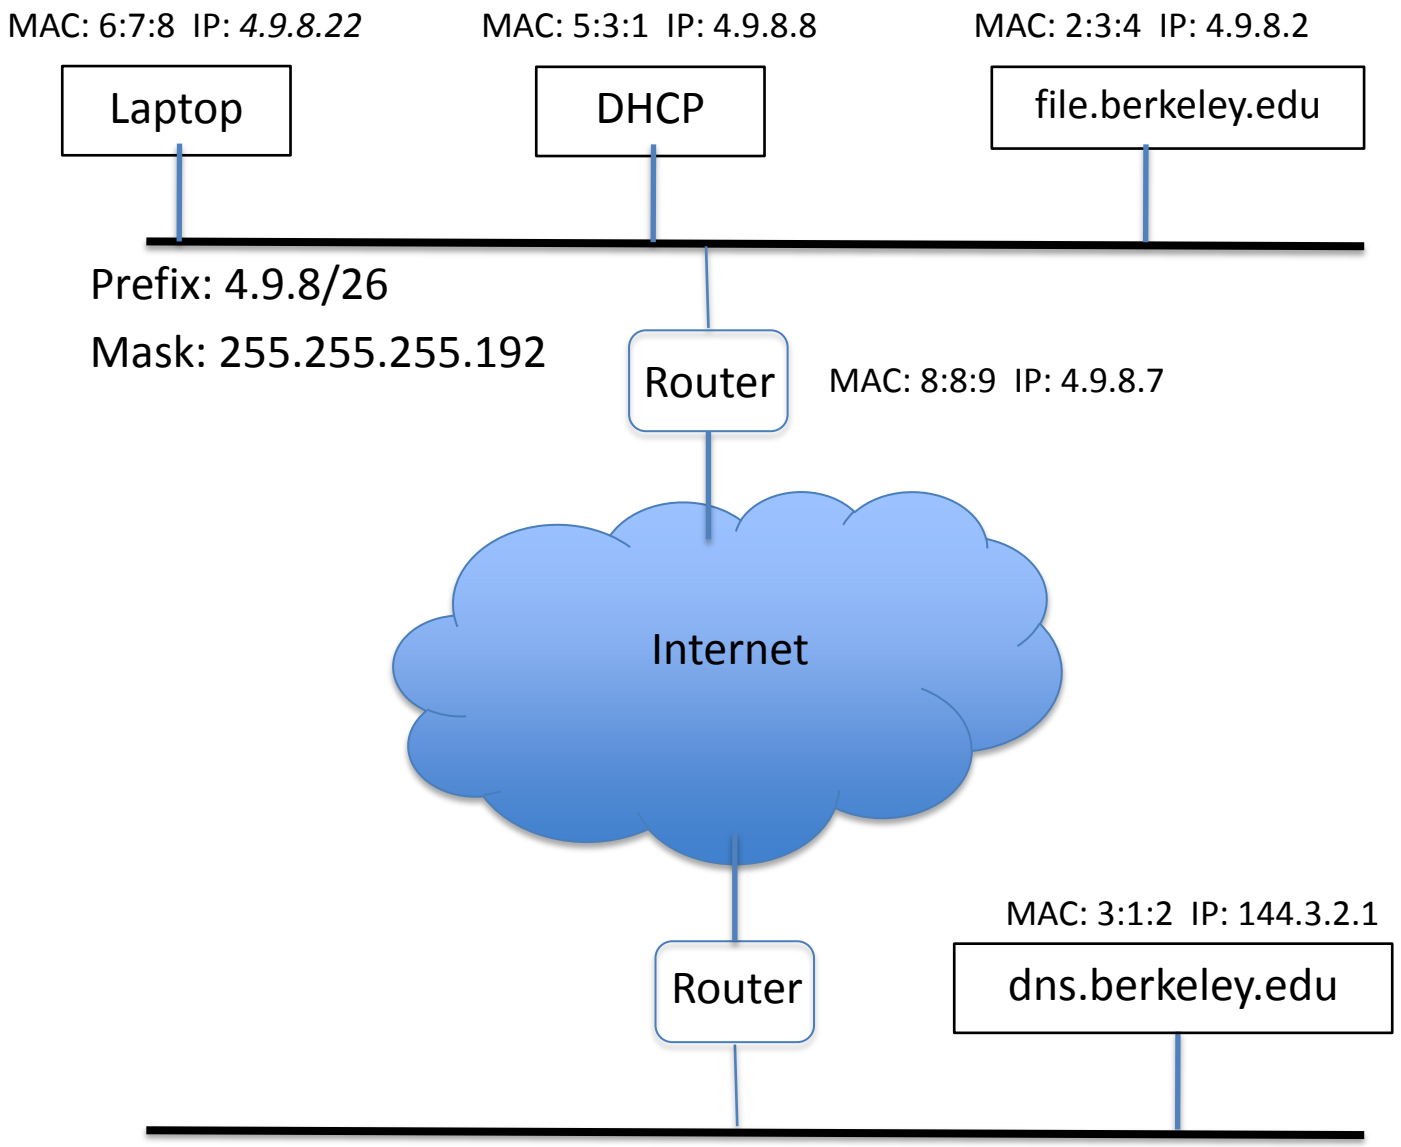


i. In the table on the following page, list the messages in the correct order and fill in the source and destination MAC and IP addresses (when applicable).

| Message | Source IP | Source MAC | Destination IP | Dest. MAC |
|---|---|---|---|---|
| Laptop: DHCP Discovery | - | 6:7:8 | 255.255.255.255 | ff:ff:ff:ff:ff |
| DHCP: DHCP Offer | 4.9.8.8 | 5:3:1 | 255.255.255.255 | ff:ff:ff:ff:ff |
| Laptop: DHCP Request (aka "DHCP Accept") | - | 6:7:8 | 255.255.255.255 | ff:ff:ff:ff:ff |
| DHCP: DHCP ACK | 4.9.8.8 | 5:3:1 | 255.255.255.255 | ff:ff:ff:ff:ff |
| Laptop: ARP Request (in order to send to DNS) | - | 6:7:8 | - | ff:ff:ff:ff:ff |
| Router: ARP Reply | - | 8:8:9 | - | 6:7:8 |
| Laptop: DNS Request | 4.9.8.22 | 6:7:8 | 144.3.2.1 | 8:8:9 |
| DNS: DNS Response | 144.3.2.1 | 8:8:9 | 4.9.8.22 | 6:7:8 |
| Laptop: ARP Request (in order to send to File) | - | 6:7:8 | - | ff:ff:ff:ff:ff |
| File: ARP Response | - | 2:3:4 | - | 6:7:8 |
| Laptop: TCP SYN to File | 4.9.8.22 | 6:7:8 | 4.9.8.2 | 2:3:4 |
| File: TCP SYNACK to Laptop | 4.9.8.2 | 2:3:4 | 4.9.8.22 | 6:7:8 |
| Laptop: TCP ACK to File | 4.9.8.22 | 6:7:8 | 4.9.8.2 | 2:3:4 |

ii. How does the laptop know the IP addresses of the DNS server and the first hop (gateway) router?

**From DHCP**

iii. How does the laptop know the file server is on the same subnet?

**It knows the netmask and it can make sure that the file server IP lies in the same subnet.**

## 2.8 Network Address Translation and FTP

When you connect to an FTP server, you are actually making two connections. First, the so-called control connection is established, over which FTP commands and their replies are transferred. Then, in order to transfer a file or a directory listing, the client sends a particular command over the control connection to establish the data connection. The data connection can be established two different ways, using active mode or passive mode.

*(a) Passive Mode*: the client sends the PASV command to the server, and the server responds with an address. The client then issues a command to transfer a file or to get a directory listing, and establishes a secondary connection to the address returned by the server.

*(b) Active Mode*: the client opens a socket on the local machine and tells its address to the server using the PORT command. Once the client issues a command to transfer a file or listing, the server will connect to the address provided by the client.

In both cases, the actual file or listing is then transferred over the data connection.

Consider a client behind a NAT box trying to use FTP using one of these two modes. Assume the NAT has not been modified to understand the semantics of FTP.

Does the client's FTP attempt succeed using Passive Mode? (yes/no) **Yes.**

Does the client's FTP attempt succeed using Active Mode? (yes/no) **No.**

## 2.9 Spanning Tree

The nodes in the graph below construct a spanning tree using the standard spanning tree protocol. The bridges (or switches), Bi, have MAC address i. After the spanning tree protocol is run, which links are in the spanning tree? **L1, L2, L5, L6, L9, L10**

i. Assume node B1 fails, and the bridges reform the spanning tree. Which links are now in the spanning tree? **L4, L7, L9, L10, L12**

iii. Assume that soon after B1 fails, B4 and B7 also fail. How many spanning trees are formed, and what are their roots? **2 spanning trees, with roots B2 and B3**

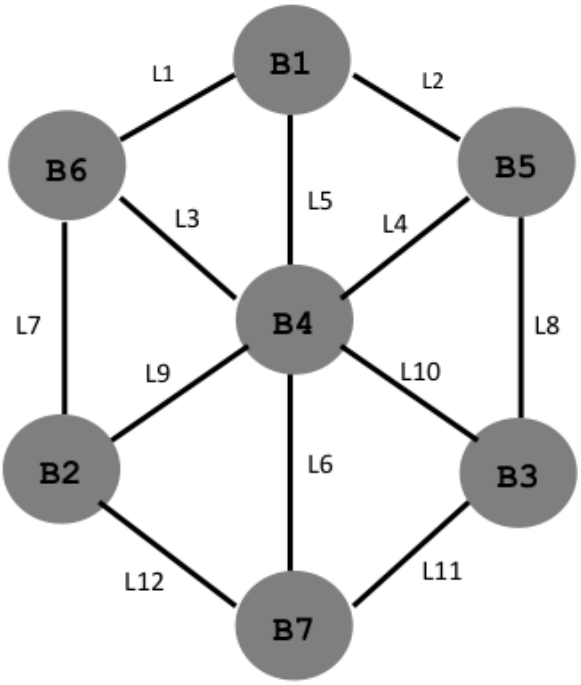


Reminder: (i) the root is the bridge with the lowest ID, and (ii) when there is more than one shortest path to the root, the path whose first hop goes through the bridge with the lower ID is chosen. Also, each link is considered to have a length of one when computing path lengths.

## 2.10 Ethernet

Hosts A and B collide on an Ethernet (in what we'll call slot 0). Each only has a single minimum sized packet to send, and no other nodes have data to send during the period in question. Which of the following three sequences (described below) are possible. Below, "slot" refers to minimum packet transmission times, and each sequence starts with the same initial collision. Assume that data consumes only a single transmission slot.

**Sequence (b)**

Sequence (a):

Slot 0: A and B both send (their first collision)
Slot 1: A and B both silent (idle)
Slot 2: A and B both silent (idle)
Slot 3: A sends, B silent (success)
Slot 4: A silent, B sends (success)

Sequence (b):

Slot 0: A and B both send (their first collision)
Slot 1: A and B both send (collision)
Slot 2: A and B both silent (idle)
Slot 3: A sends, B silent (success)
Slot 4: A and B both silent (idle)
Slot 5: A silent, B sends (success)

Sequence (c):

Slot 0: A and B both send (their first collision)
Slot 1: A sends, B silent (success)
Slot 2: A and B both silent (idle)
Slot 3: A and B both silent (idle)
Slot 4: A silent, B sends (success)

## 2.11 Multicast Routing

Consider the network below, with each link having the same cost (in the metric used in unicast routing). The members of the multicast group are denoted by m1, m2, m3, and routers in the network are denoted by A, B, C, D, E, F.

(i) Assume that CBT multicast routing is used, with the center (core) located as marked in the network. Describe network paths by the series of routers the packets pass through.

What path do packets sent by M1 to the group take to reach M2? **M1-A-B-C-E-M2**

What path do packets sent by M3 to the group take to reach M2? **M3-D-C-E-M2**

(ii) Assume that single-source multicast trees are used, and ignore the node labeled as core.

What path do packets sent by M1 to the group take to reach M2? **M1-A-F-E-M2**

What path do packets sent by M3 to the group take to reach M2? **M3-D-E-M2**

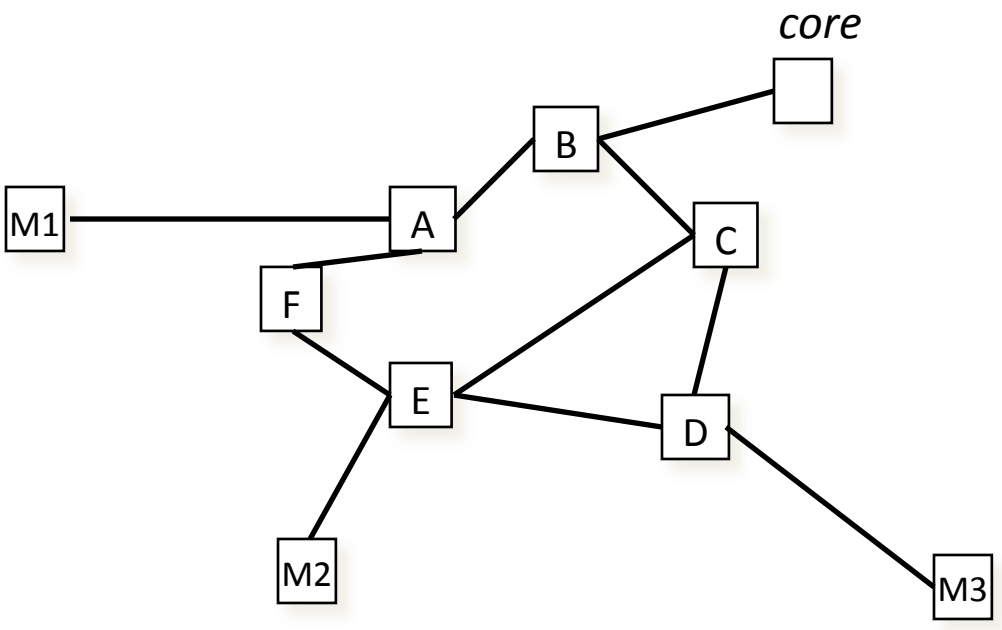

## 3.5 BGP

i) Consider the network in Figure 1 (on the next page), with ASes A, B, C, D. Each AS has some number of routers (labeled as A1, A2, etc.) and the domains are connected internally and with each other by the links depicted in the figure. Assume that eBGP and iBGP are used for interdomain routing, and that ASes A and D are using RIP for intradomain routing while ASes B and C are using OSPF for intradomain routing. Prefix x hangs off an interface on router C3. For the following answer, use one of these options:

a) OSPF
b) RIP
c) eBGP
d) iBGP

Router D3 learns about prefix x from which routing protocol?__ **c** __

Router D1 learns about prefix x from which routing protocol? __ **d**__

Router A3 learns about prefix x from which routing protocol? __ **c**__

Router A1 learns how to reach router A3 from which routing protocol? __ **b**__

Will router A1 use interface 1 or interface 2 to reach prefix x? _ **1**__

ii) Consider an interdomain network with domains A through F. For simplicity, assume that destinations in this problem are domains, not prefixes. Recall that routes are expressed in terms of the series of domains: e.g., [A-B-C] denotes a route that started with domain A and went to domain B and then to domain C (which is the destination). Domains always advertise the route to themselves (i.e., domain X advertises the path [X] to all peers, customers, and providers). The following connectivity/business relationships exist:

- B is a customer of A
- C is a customer of A
- D is a customer of B
- E is a customer of B
- F is a customer of C
- G is a customer of C
- B and C are peers
- E and F are peers

Assuming that each domain's routing policies follow normal business practice, and that BGP has converged,

- What routes does A advertise to B? _ **[A] [A-C] [A-C-F] [A-C-G]**
- What routes does C advertise to B? _**[C] [C-F] [C-G]**
- What routes does E advertise to B? _**[E]**

- What routes does F advertise to E? ___**[F]**___
- What path do packets from E take to F? __**E-F**__
- What path do packets from D take to F? ___**D-B-C-F**___
- What path do packets from D take to G? ___**D-B-C-G**

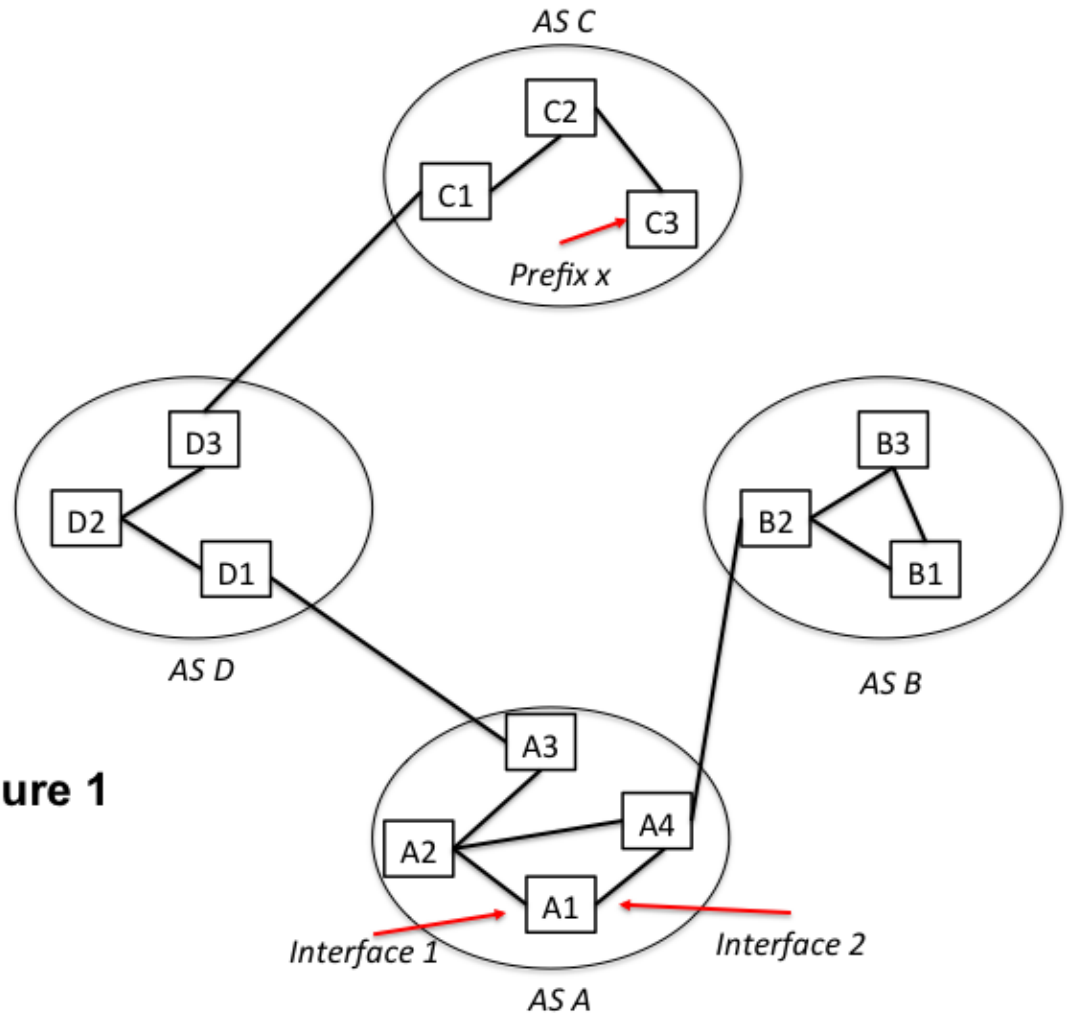


**Figure 1**