Module 4: Local Area Networks (LANs); VLANs and Networking Devices

Dr. Natarajan Meghanathan
Professor of Computer Science
Jackson State University
Jackson, MS 39217
Phone: 601-979-3661
E-mail: natarajan.meghanathan@jsums.edu

Module 4 Topics

- 4.1 Local Area Networks (LAN)
  - LAN Topologies, Ethernet and Wireless LANs
- 4.2 Networking Devices to Extend LANs
  - Repeater, Hub, Bridge, Switch
- 4.3 Virtual LANs

4.1 Local Area Networks (LANs)

Local Area Networks (LANs)

- A LAN consists of a single shared medium to which many computers attach.
- The computers “coordinate” to access the medium and transmit packets.

- LAN eliminates duplication of messages when a computer wants to send a message to more than one computer simultaneously.

- LANs are used mainly for local communications – For long distance communication, the time spent to co-ordinate use of the shared medium becomes significantly higher than the time required to send data. Note that the time required to communicate depends on the distance.

- Point-to-point connections are preferred for long-distance communication.

LAN Topologies

- Bus Topology
  - All Computers are connected to a single cable.
  - Any computer can send data to any other computer. Computers need to co-ordinate their transmissions to avoid packet collisions.
  - A signal sent by one computer travels down the cable and all the computers can receive the signal.
- Ring Topology
  - Computers are connected in a closed loop.
  - Data sent by a computer is passed from one computer to the next in the loop until the data reaches the destination computer/ sender.
- Star Topology
  - Each computer is attached to a central point called a hub.
  - The hub accepts data from the sending computer, forwards to everyone else (thereby, the message reaches the destination computer).

LAN Topologies

- Bus Topology

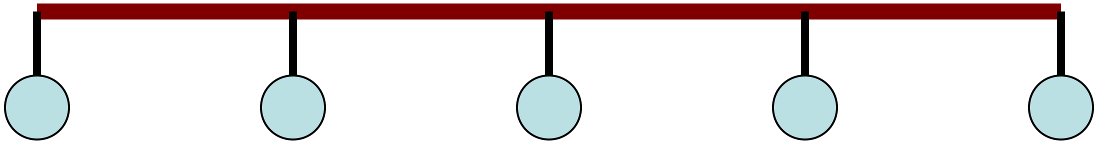

- Ring Topology

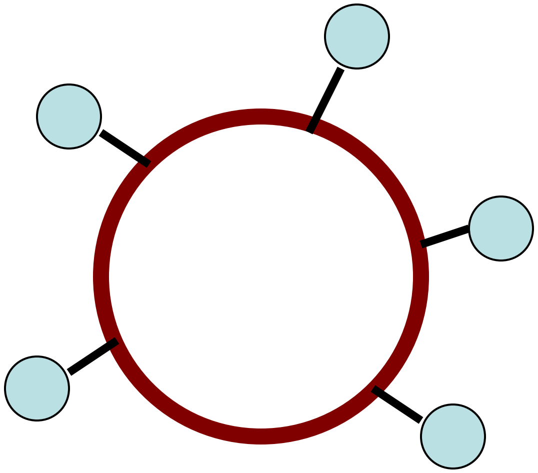

- Star Topology

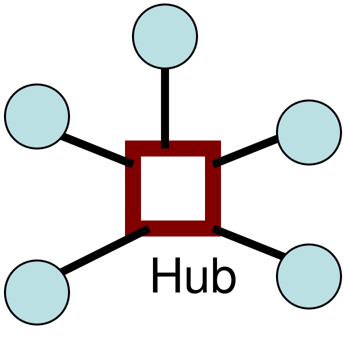

Comparison of LAN Topologies

- Physical topology indicates how computers are connected to each other; whereas, logical topology illustrates how the computers communicate with each other.
- The bus topology and star topologies are physically different; but, logically equivalent to each other as a message sent from one computer is forwarded by the hub to all the other computers connected to it.
- The bus topology can endure the failure of any computer; but, gets disconnected with any break in the cable.
- The ring topology is susceptible to failure with break in the cable; but, it could be made more fault tolerant with the use of a backup ring.
  - Ring topologies typically use a token to grant access. There is only token in the ring at any point of time and the machine that gets the token is the only machine that can access the ring. If the machine crashes while holding the token, the token is lost and has to be regenerated.
- The star topology is also susceptible to a single point of failure (hub); but, can survive the failure of any individual computers and/or any links.

Ethernet

- Uses the bus topology
- Uses the Manchester Encoding standard for physical layer communication
  - Employs a 64-bit preamble (alternating 0s and 1s) to precede a frame and this is used to synchronize the source and the destination
- The length of a single Ethernet segment can be up to 500m; while the minimum separation between two computers attached to the Ethernet must be 3m.

Type	**Cable**	**Bandwidth**
10Base5	Coaxial	10Mbps
10BaseT	Shielded Twisted Pair	10Mbps
10BaseF	Fiber optic	10Mbps
Fast Ethernet	Twisted pair, fiber optic	100 Mbps
Gigabit Ethernet	Twisted pair, fiber optic	1000 Mbps

Ethernet – CSMA (Carrier Sense Multiple Access)

- No centralized controller to co-ordinate frame transmission
- All the hosts attached to an Ethernet network participate in a distribution co-ordination scheme called CSMA.
- Idea: Use the electrical activity in the cable to determine the status of the cable (whether a transmission is in progress or not?)
- **Carrier** – the electrical signal
- **Carrier sense** – checking the medium whether any electrical signal is in transmission
- **CSMA** – determining whether to transmit or not based on the result of carrier sense
- When a potential sender does not sense the presence of any electrical signal in the medium, it takes an independent (individual) decision to transmit a frame.
- A host wishing to transmit should check the medium for any on-going transmission. If any electrical signals are detected, the host considers it as a failure to transmit the frame and attempts to retransmit the frame.

Collision Detection

- Since a host decides to transmit when it senses no signal in the medium, CSMA is prone to collision (interference) of signals.
- Once collision occurs, both the frames get garbled and the signals reaching the hosts are not useful.
- To detect collision, a sending station must ensure that it listens (transmits the signal) to the channel unless it can make sure that the first bit of the data has reached the receiver.
- During this time, if the sender detects that the signal in the medium is different from what it is transmitting (happens when there is a collision), the sender decides a collision has occurred and stops transmitting.
- After a collision, the signal transmitted in the Ethernet network is generally a garbled frame and all hosts will be able to sense and distinguish a garbled signal from a valid frame signal.

Determining the Minimum Frame Size

- How to achieve collision detection? Solution: For a given Ethernet network length and bandwidth, Ethernet frames require a minimum frame size.
- Propagation delay, d = $\dfrac{\text{distance traveled by the signal}}{\text{velocity of the signal in the medium}}$
- Transmission delay = $\dfrac{\text{Packet size (bits)}}{\text{Bandwidth of the medium (bps)}}$
- Round-trip time, RTT= 2*Propagation delay
- Worst-case Scenario: Consider a collision between frames transmitted by the hosts on the two ends of the network

Determining the Minimum Frame Size

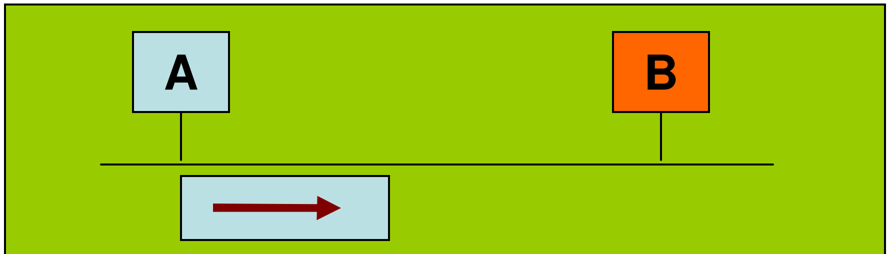

At time t: A transmitting a frame to B

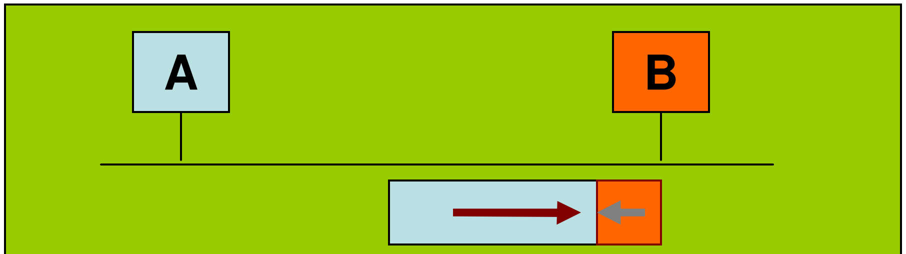

Just before t+d: B begins transmitting

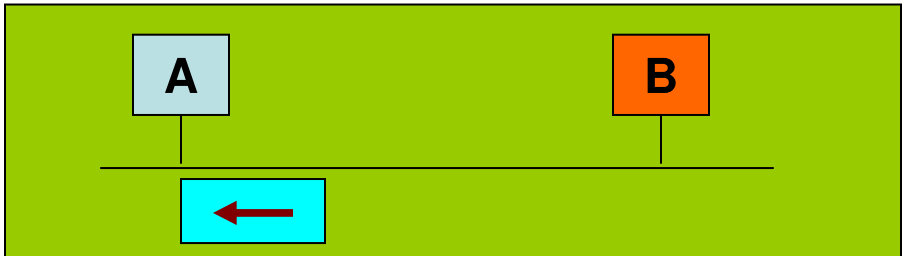

At time t + 2d: A sees garbled frame

A needs to transmit for duration 2d, i.e., RTT
In other words, transmission delay = RTT

$$\frac{\text{Packet size (bits)}}{\text{Bandwidth of the medium (bps)}} = \frac{\text{2 * Length of the Ethernet network}}{\text{Velocity of light in the medium}}$$

How to Recover from Collision?

- If two hosts whose frames collided, wait for the same amount of time, before attempting retransmission of the frame, we will only have a sequence of collisions and nothing else.

- A third host that attempted to transmit its frame and ended up sensing a garbled signal, resulting from the collision of frames of two other computers, should also refrain itself from transmitting and consider its transmission attempt as a failure and try for retransmission.

- We need the hosts to wait for a random amount of time during each collision and the probability that the amount of time a host waits equals the waiting time of another host should be very very … less.

How to Recover from Collision?

- **Binary Exponential Backoff algorithm**
- Before any collision occurred for the frame to be transmitted, let d be the maximum delay that a host is configured to wait.
- For the first retransmission attempt, the host waits for a time randomly selected between 0…d and then attempts to retransmit. If the medium is busy or a collision occurs, the host waits for a second retransmission.
- For the second retransmission attempt, the host waits for a time randomly selected between 0…2d
- For the third retransmission attempt, the host waits for a time randomly selected between 0…4d
- For the fourth retransmission attempt, the host waits for a time randomly selected between 0…8d
- In general, if k is the number of retransmission attempts, the host chooses a wait time randomly from 0 … $(2^k-1)*d$

- If the host interface cannot successfully transmit the packet after the maximum retransmission attempts (=16), then it reports a medium access failure.

Sample Question: Ethernet Frame Length

- Consider an Ethernet segment operated at 100 Mbps, and a segment was limited to 500meters in length. Assume the signal propagates down the cable at two-thirds the speed of light. Compute the minimum Ethernet frame size?

Min. Transmission Delay = 2 * Propagation Delay

$$\frac{\text{Min. Ethernet Frame Size}}{\text{Channel Bandwidth}} = 2 * \frac{\text{Channel Length}}{\text{Speed of the signal on the channel}}$$

$$\text{Min. Ethernet Frame Size} = \frac{2 * 500\ \text{m} * 100 * 10^{6}\ \text{bits/sec}}{0.66 * 3 * 10^{8}\ \text{m/sec}}$$

= 500 bits

Ethernet Frame Format

- An Ethernet frame consists of a fixed-length header, a variable-length payload, and a fixed-length Error detection code (CRC)
- If the frame size along with the data falls below the minimum frame size required under the hardware technology, the data is extended with zeroes (called as “padding “).

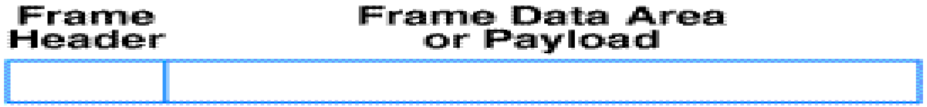

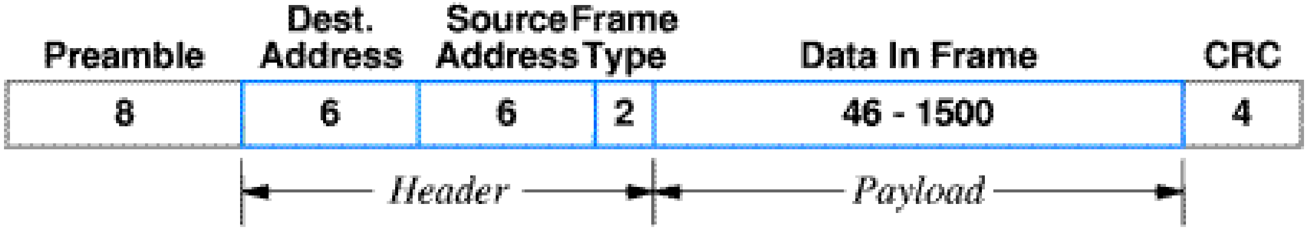

Frame Type – identifies the protocol in the Network Layer to which the Payload (data) needs to be handed over.

802.11 Wireless LANs

- Wireless LAN hardware use antenna to broadcast RF signals (through the air), which are received by antenna attached with other computers.

- The antennas of all the computers are configured to transmit at the same radio frequency. Hence all the computers connected to a wireless LAN need to co-ordinate to access transmission.

- Wireless LAN transmitters use low power; hence, transmissions propagate only to a limited radius around the sender. The distance beyond which the transmissions of a wireless device do not propagate significant enough to cause interference is called the transmission range.

Hidden Terminal Problem

- Lack of full connectivity of the transmitters in a wireless LAN leads to the hidden terminal problem.
- **Hidden terminal problem when using CSMA/CD:**
- There exists a common receiver B that is in the transmission range of two senders A and C.
- A and C are not in the transmission range of each other. So, when A or C attempt to transmit to B, they sense the medium to be free of transmissions and hence send their frame to B.
- Simultaneous transmissions or transmissions closely spaced in time from both A and C will have their frames colliding at B.

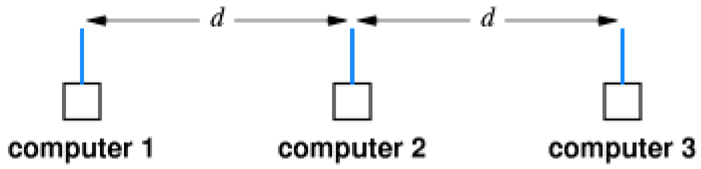

CSMA with Collision Avoidance (CA)

- Consider host A attempting to transmit to B.
- Even, if the medium surrounding A is free of any ongoing transmission, A wishes to make sure that the medium surrounding B is also free of any transmission. In other words, the sender would like to reserve the medium surrounding the transmission range of the receiver.
- A sends a Request to Send (RTS) frame to B. B on receiving the RTS, responds with a Clear to Send (CTS) frame if it is free to receive any transmission.
- The neighbors of B on receiving the CTS frame will defer their intended transmissions and wait for a time equivalent to a transmission of a packet.
- A on receiving the CTS frame, transmits the data frame to B.

- What if the RTS packet of one host collides with the RTS of another? Solution: Apply random backoff similar to that of Ethernet. Since the size of a RTS frame is much shorter than the size of data frames, the probability of collisions of RTS frames is far less compared to the collision of data frames.

CSMA/ CA

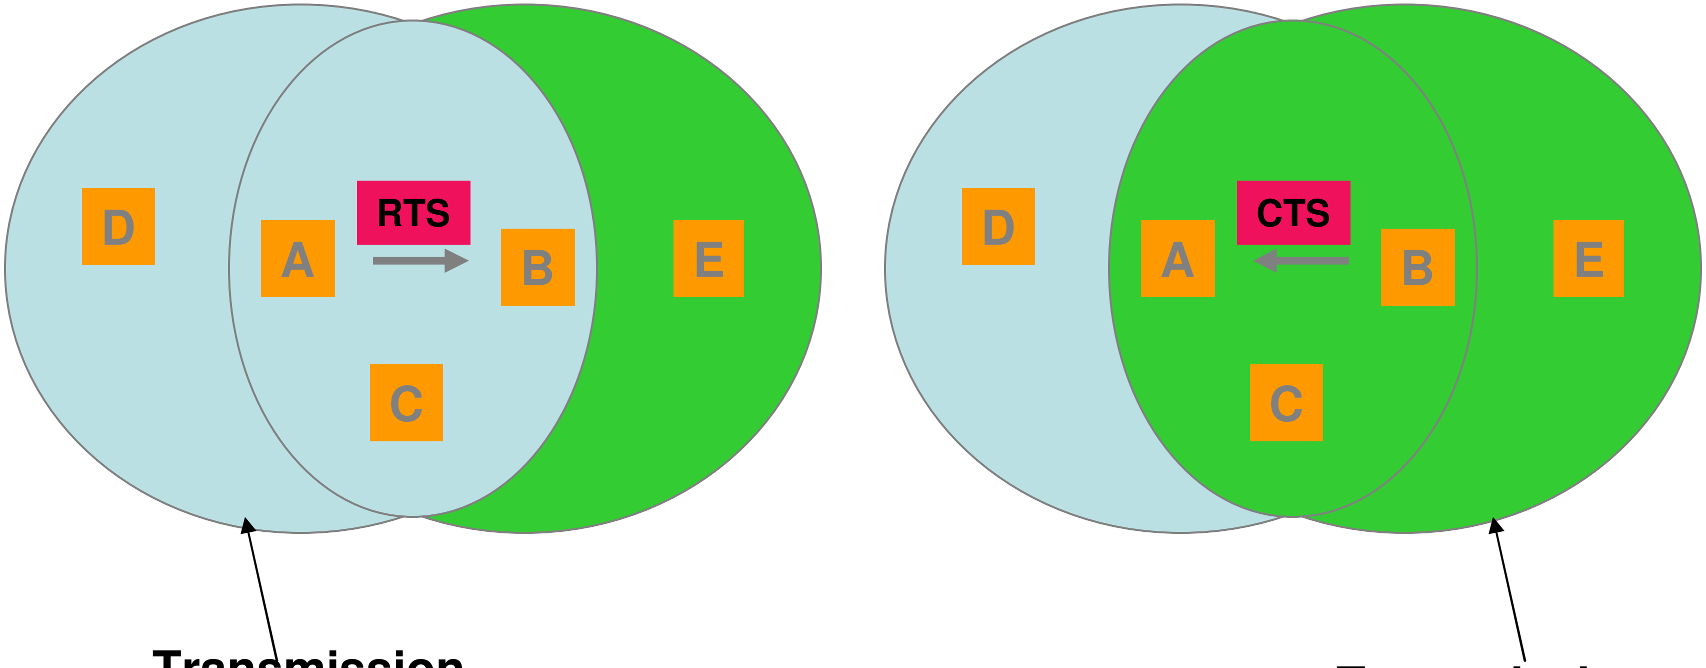

Transmission Range of A

Transmission Range of B

4.2 Networking Devices to Extend LANs

Motivation for Extended LANs

- How to facilitate interactions between people who are in the same organization, but separated with distances larger than a few hundred meters?
  - Electrical signals do not travel for long distances without appreciable decrease in signal quality.
- As the length of the LAN increases, more computers are added to the LAN and the fairness of access to the computers is tough to be guaranteed.
  - The delay incurred to gain control of the medium increases proportional to the number of computers attached to the LAN and needs to be bounded.
- Solution: Extend LANs using devices like repeaters, hubs (both layer 1), bridges, switches (both layer 2) and routers (layer 3)
  - **With repeaters and hubs, the LANs are in the same collision domain as well as the same broadcast domain**
  - **With bridges and switches, the LANs are in separate collision domains, but the same broadcast domain**
  - **With a router (or layer-3 switches), the LANs are in separate collision domains as well as separate broadcast domains.**

Repeaters

- A repeater connects two cable segments, amplifies and sends all electrical signals received from one segment to the other segment without waiting for the complete frame to arrive.
- A repeater does not have a physical address; it directly attaches to the Ethernet cables on both ends.

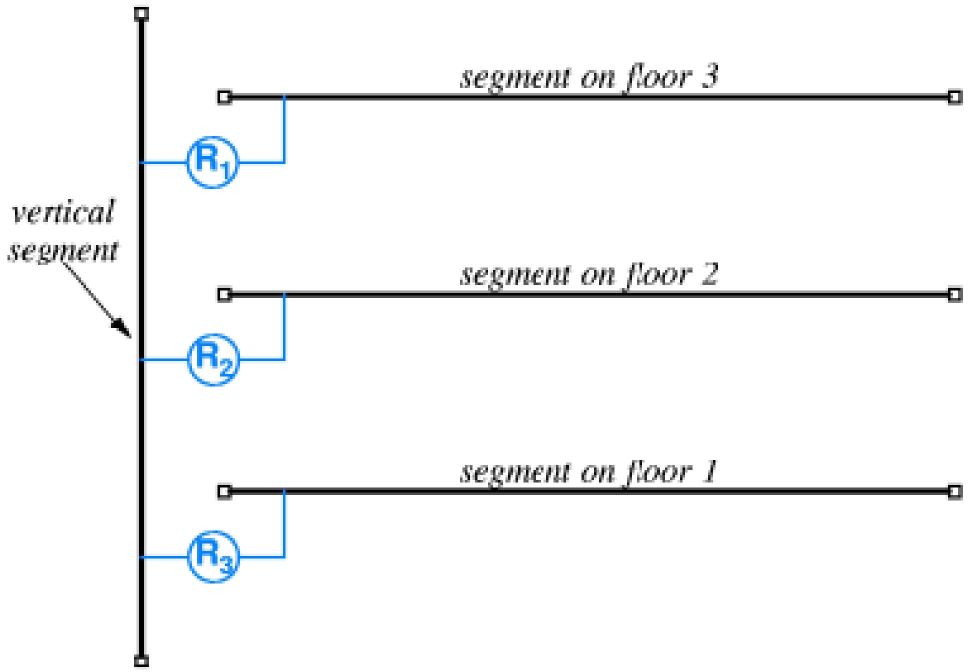

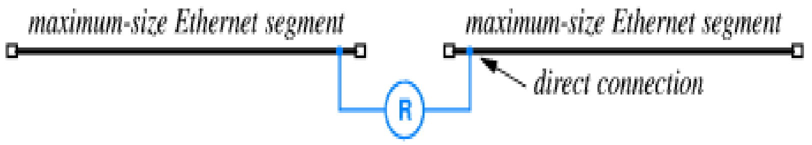

- **The maximum number of repeaters between any two hosts in an Ethernet segment is 4.**

- **A repeater propagates a copy of all electrical signals (including collision, interference) received in one segment onto another.**

Hubs

- A hub is a multi-port repeater, with one machine or LAN connected to each of its ports.
- A hub forwards the message received from one port to all the other ports.
- A hub does not have a physical address.
- A hub is typically used to implement a 10Base-T Ethernet in the form of a star topology.

Bridges

- A bridge connects two LAN segments; forwards only complete, correct frames from one segment to another.
- A bridge splits collision domains, but not broadcast domains
- The bridge uses the same network interface as to that of a conventional computer and listens to the segments it attaches in promiscuous mode (listening to on-going traffic in the segment).

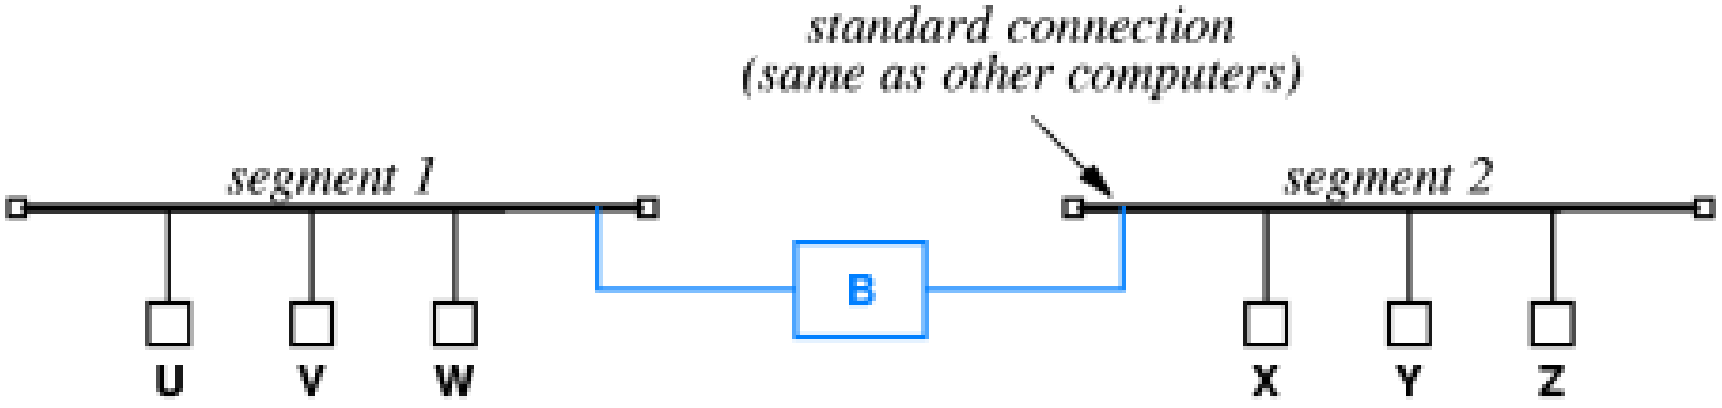

Learning Bridges (Adaptive Bridges)

- A bridge listens in promiscuous mode, on the segment to which it attaches and forms a list of computers attached to the segment.

Learning process:

- First, the bridge extracts the physical source address from the frame header and adds it to the list of computers attached to the segment.
- Second, the bridge extracts the physical destination address from the frame header and uses the address to determine whether to forward the frame.

- **Steady state** – a state in which a bridge forwards a frame only if necessary.
- Startup state – A state in which the bridge learns the location of computers. Frames for which the location of the destination is not known, are forwarded to all the segments other than the one on which the frame arrived.

- Each computer sends at least one frame after booting up, to help the bridge reach its steady state quickly.

Learning Bridges (Adaptive Bridges)

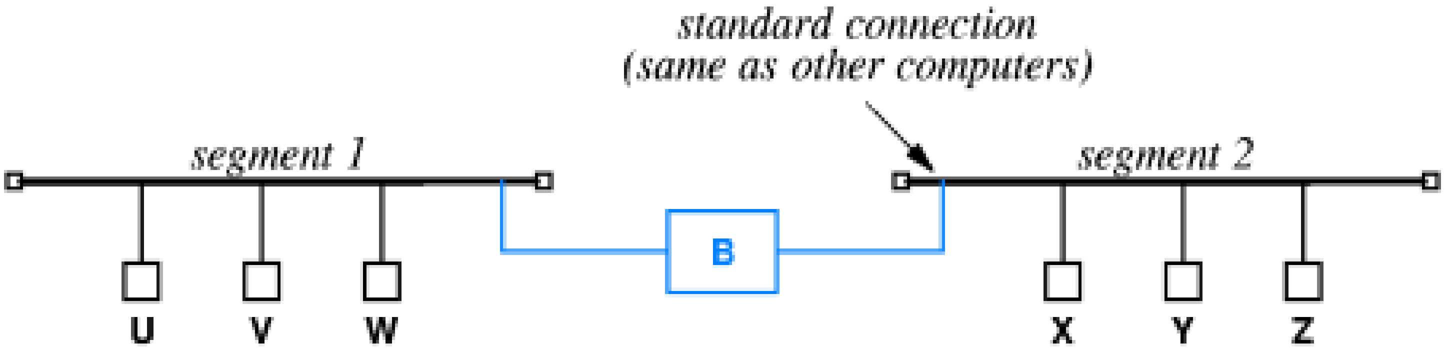

Event	Segment 1 List	Segment 2 List
Bridge boots	–	–
U sends to V	U	–
V sends to U	U, V	–
Z broadcasts	U, V	Z
Y sends to V	U, V	Z, Y
Y sends to X	U, V	Z, Y
X sends to W	U, V	Z, Y, X
W sends to Z	U, V, W	Z, Y, X

A Cycle of Bridges

- Consider what happens when a broadcast frame is sent from segment a?
- Segment ‘a’ sends to ‘b’ and ‘c’; both ‘b’ and ‘c’ send to ‘d’ separately. Segment ‘d’ again forwards the frame received from ‘b’ to ‘c’ and the frame received from ‘c’ to ‘b’. This will continue for ever, until one bridge stops forwarding.

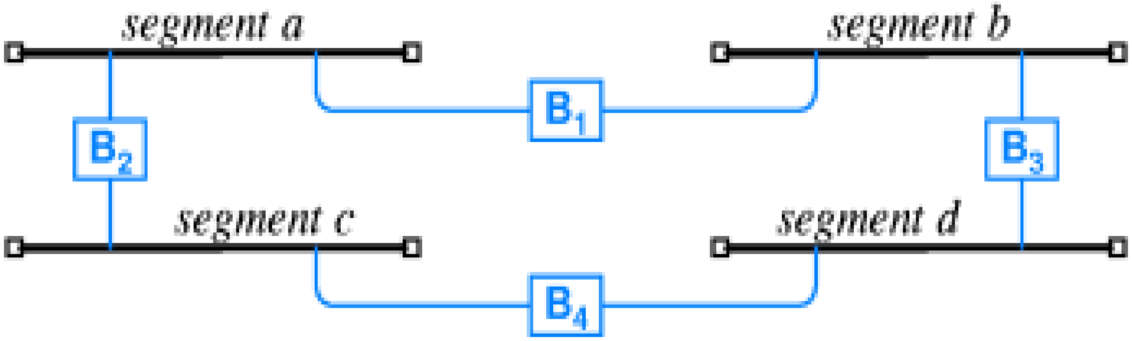

- To prevent infinite loops, two conditions should not occur simultaneously in a bridged network: (i) all bridges forward all frames (ii) the bridged network contains a cycle of bridged segments.

Spanning Tree

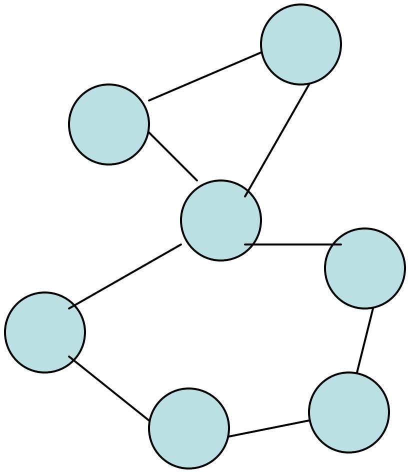

A Cyclic Graph, G

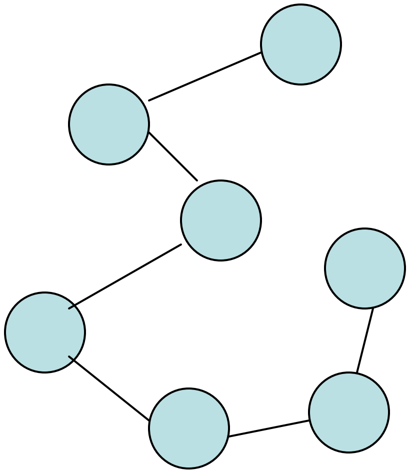

A Spanning Tree for G

In a tree, there can be exactly one path between any two nodes.

If the tree includes all the nodes in the original graph, then the tree is called a spanning tree of the graph.

Distributed Spanning Tree (DST)

- **Idea:** For each bridge, select the interfaces (ports) over which they will forward frames (designated port) or will communicate with the root as well as receive frames (root port) or will not communicate (blocked port: and a cycle can be avoided).
- **Assumption:** Each bridge has a unique identifier.
- **Root bridge:** The bridge with the smallest ID.
- **We decide the roles for the ports of the bridges in the following order.**
- **Root port (RP) for a bridge: (other than the root bridge)** The port that is on the shortest path to the root bridge and is used by the bridge to communicate with the root bridge as well as receive packets that are seen in the LAN
  - Tie, choose the port leading to the neighbor bridge with the smallest ID
- **Designated port (DP) for a bridge:** The port that is used to forward a message to the attached LAN.
  - A port incident on a bridge becomes a designated port if the other end of the LAN/link is a root port.
  - If a LAN/link has no bridges whose incident port is a root port, the port for the bridge with the lowest ID becomes the designated port.
- **Blocked port (BP) for a bridge:** The port that is NOT used to send or receive messages.
  - A port that has not been classified either as a root port or a designated port.

DST Rule of Thumb

- Every bridge (other than the root bridge) should have a root port (which is on the shortest path to the root bridge)
  - At each non-root bridge, messages seen in the LAN (incl. ones from the root bridge) are received through the root port and forwarded further if a port on the bridge is also a designated port.

- Every LAN should have ONLY ONE among its incident bridge ports chosen as the designated port (only through which messages can be forwarded)
- If a link/LAN has one end has a bridge with a root port, the other end of the link/LAN should be classified as a bridge with a designated port.
- If a link/LAN does not have bridges with root port or designated port yet setup, then the port of the bridge with the smallest ID becomes the designated bridge.

- If a port fails to get classified as the root port or designated port, it becomes a blocked port.

Example 1: DST

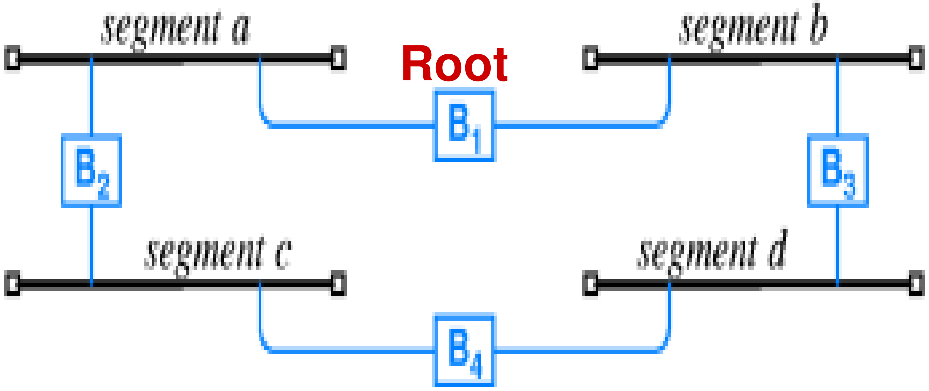

Final DST

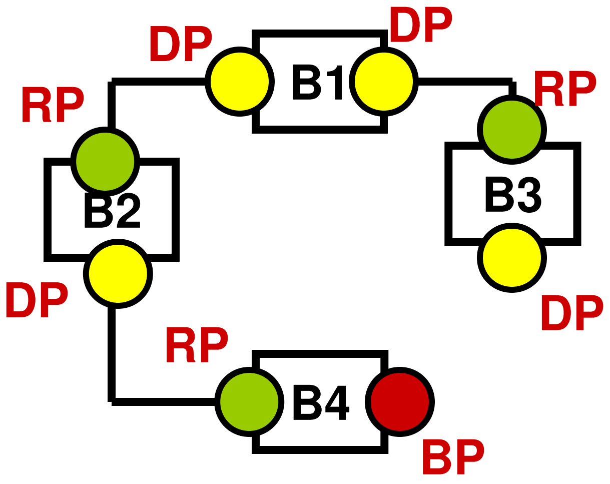

Step 1: Choosing Root Ports

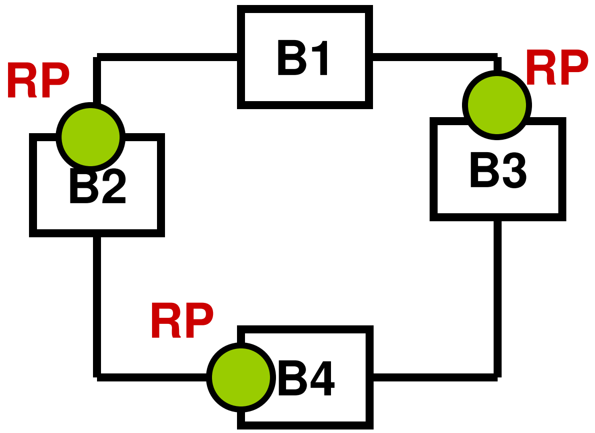

Step 2: Choosing Designated Ports

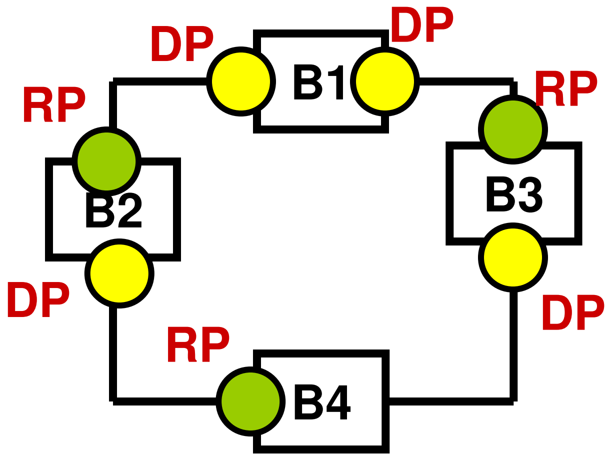

Step 3: Choosing Blocked Ports

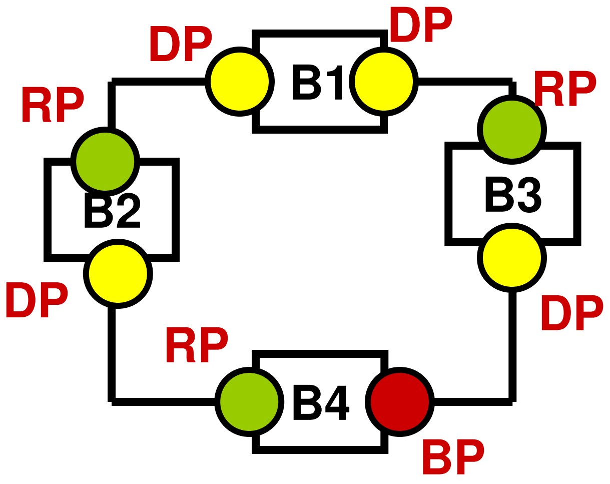

Example 1: DST

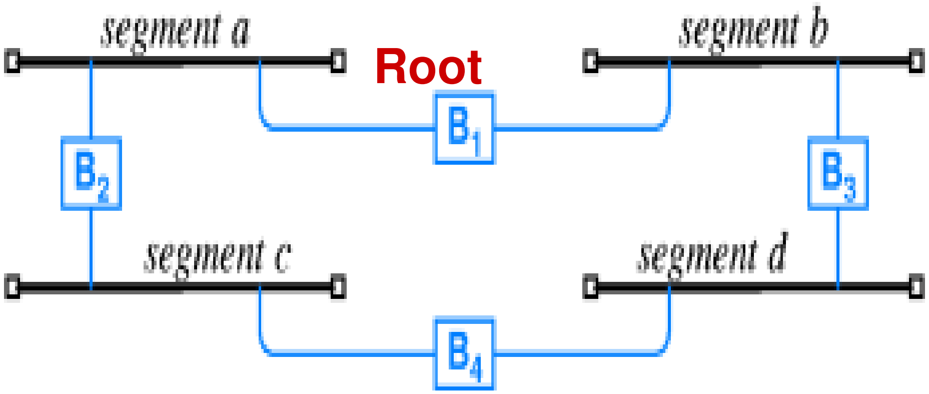

Final DST

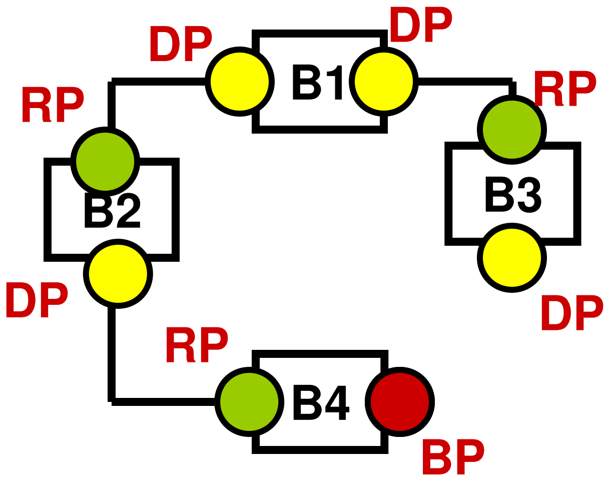

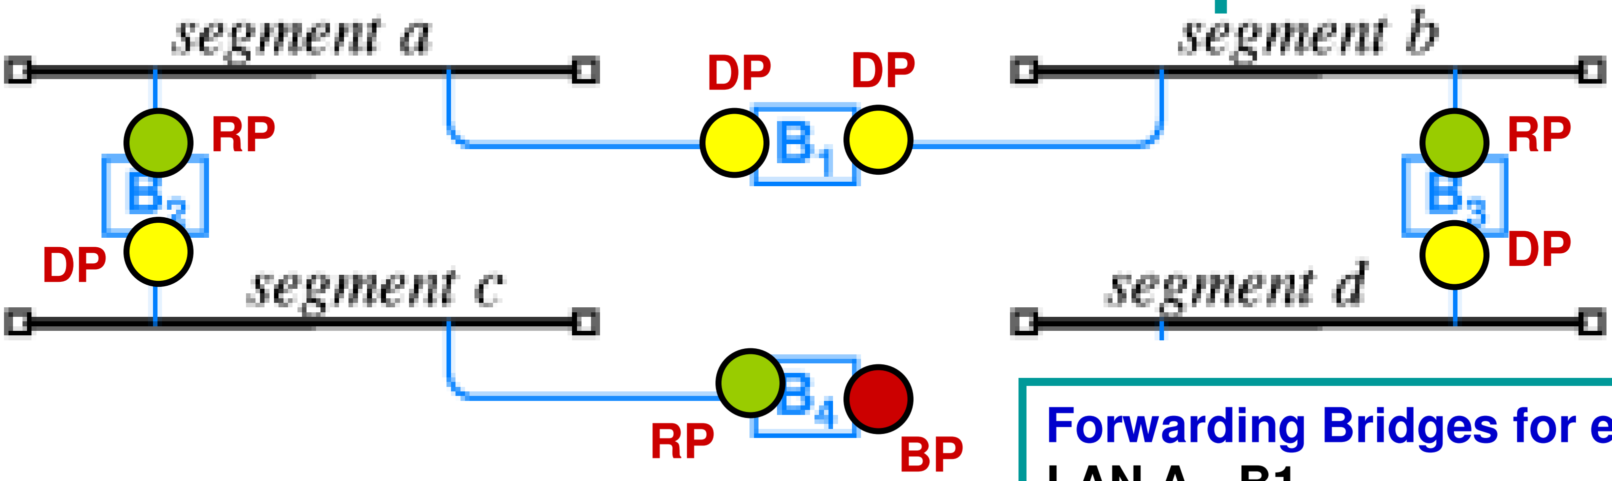

Forwarding Bridges for each LAN

LAN A	B1
LAN B	B1
LAN C	B2
LAN D	B3

Example 2: DST

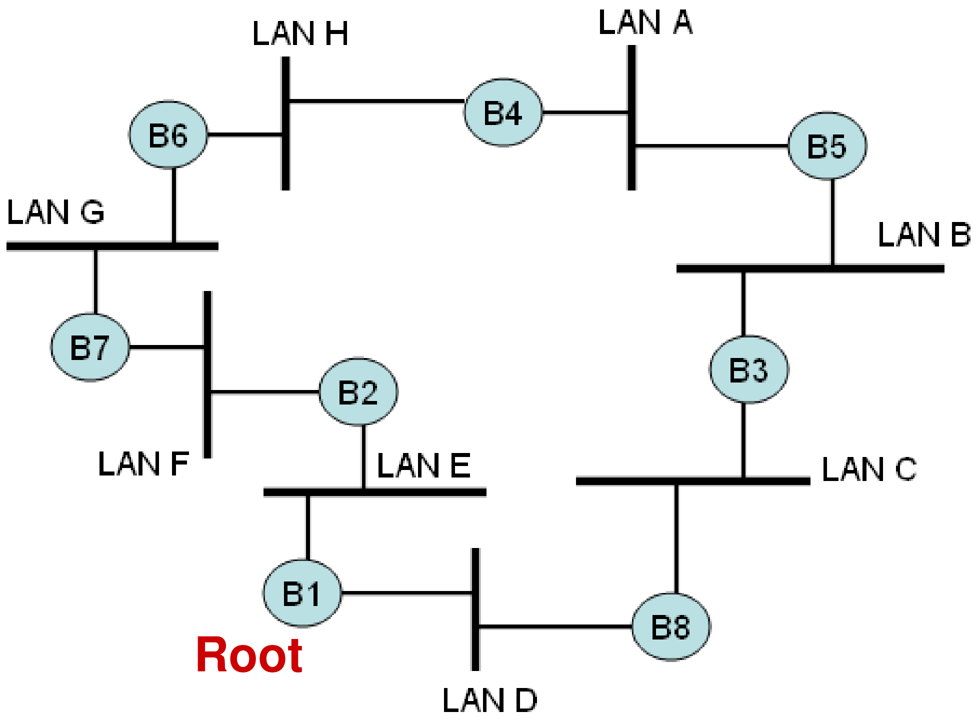

Step 1: Choosing Root Ports

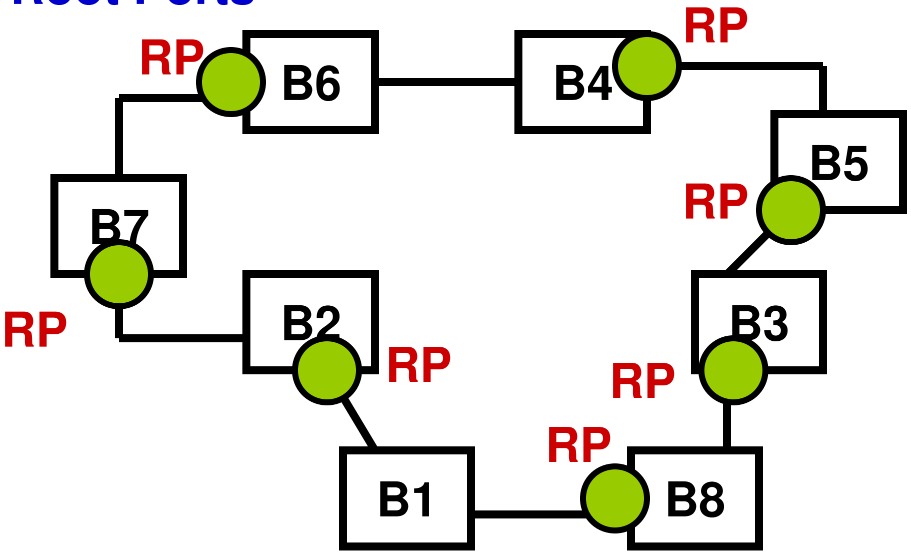

Step 2: Choosing Designated Ports

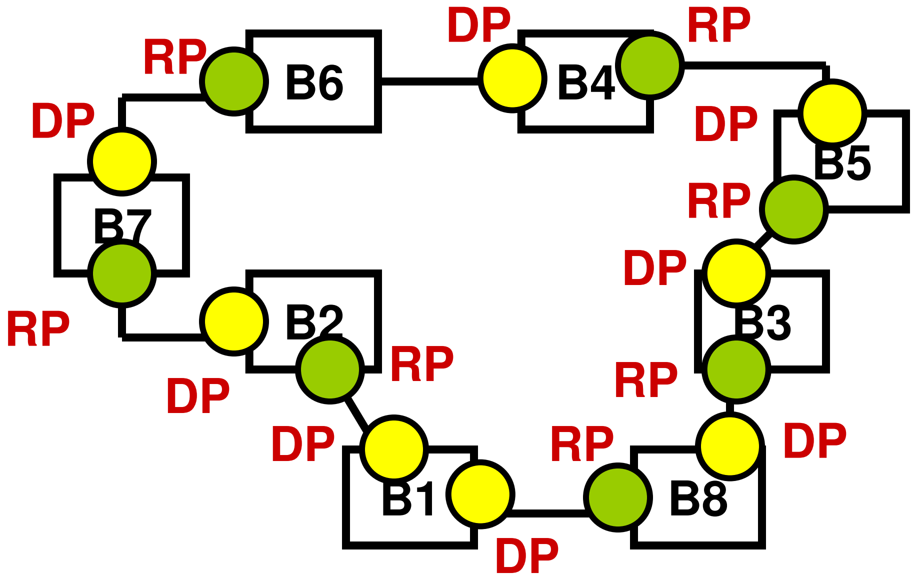

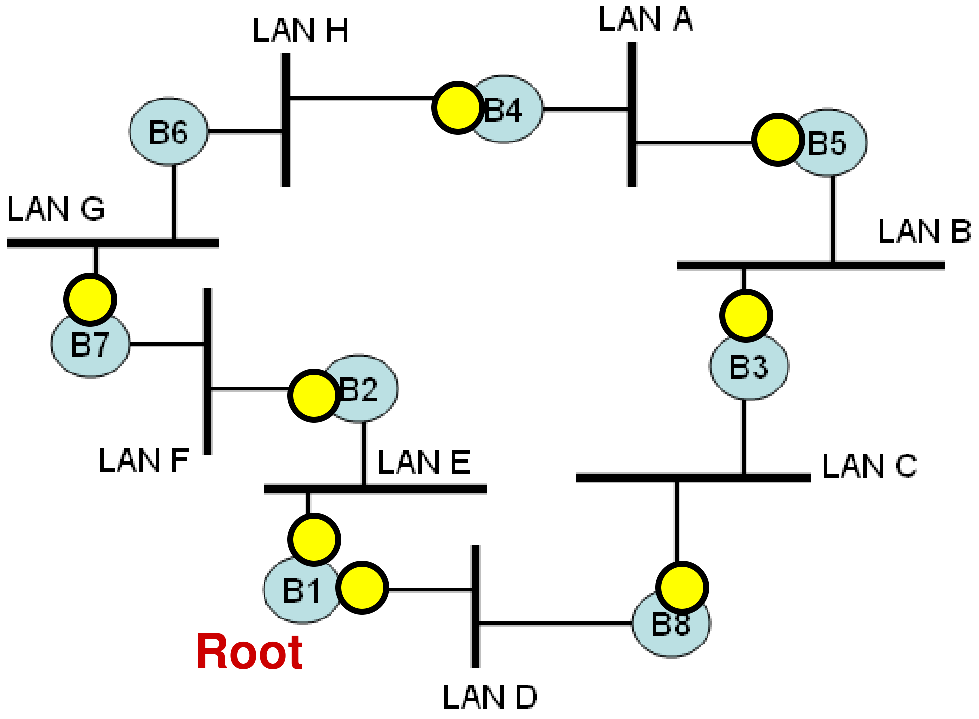

Forwarding Bridges for each LAN	
LAN A	B5
LAN B	B3
LAN C	B8
LAN D	B1
LAN E	B1
LAN F	B2
LAN G	B7
LAN H	B4

Step 3: Choosing Blocked Ports

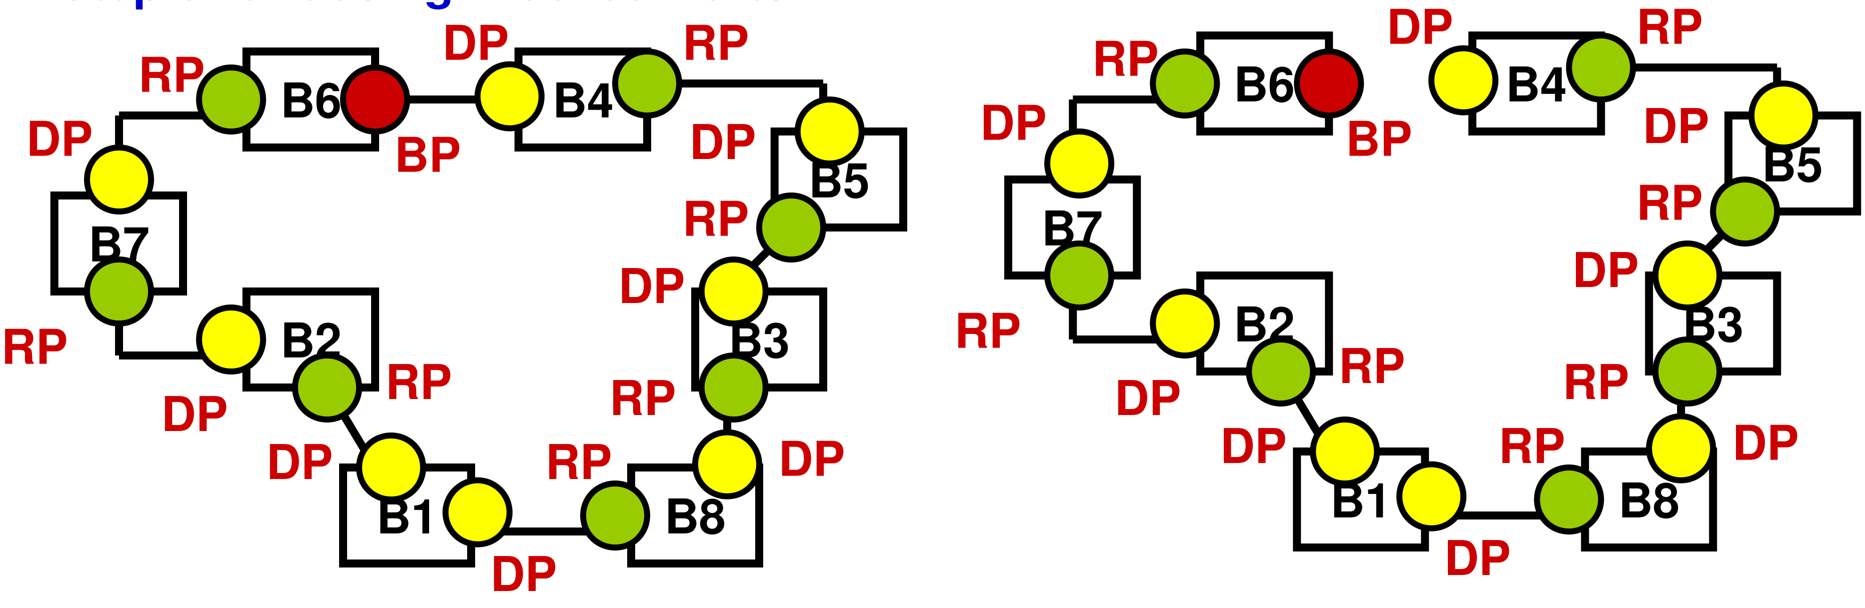

Example 3: DST

Note: This example is applicable
for switches: multi-port bridges

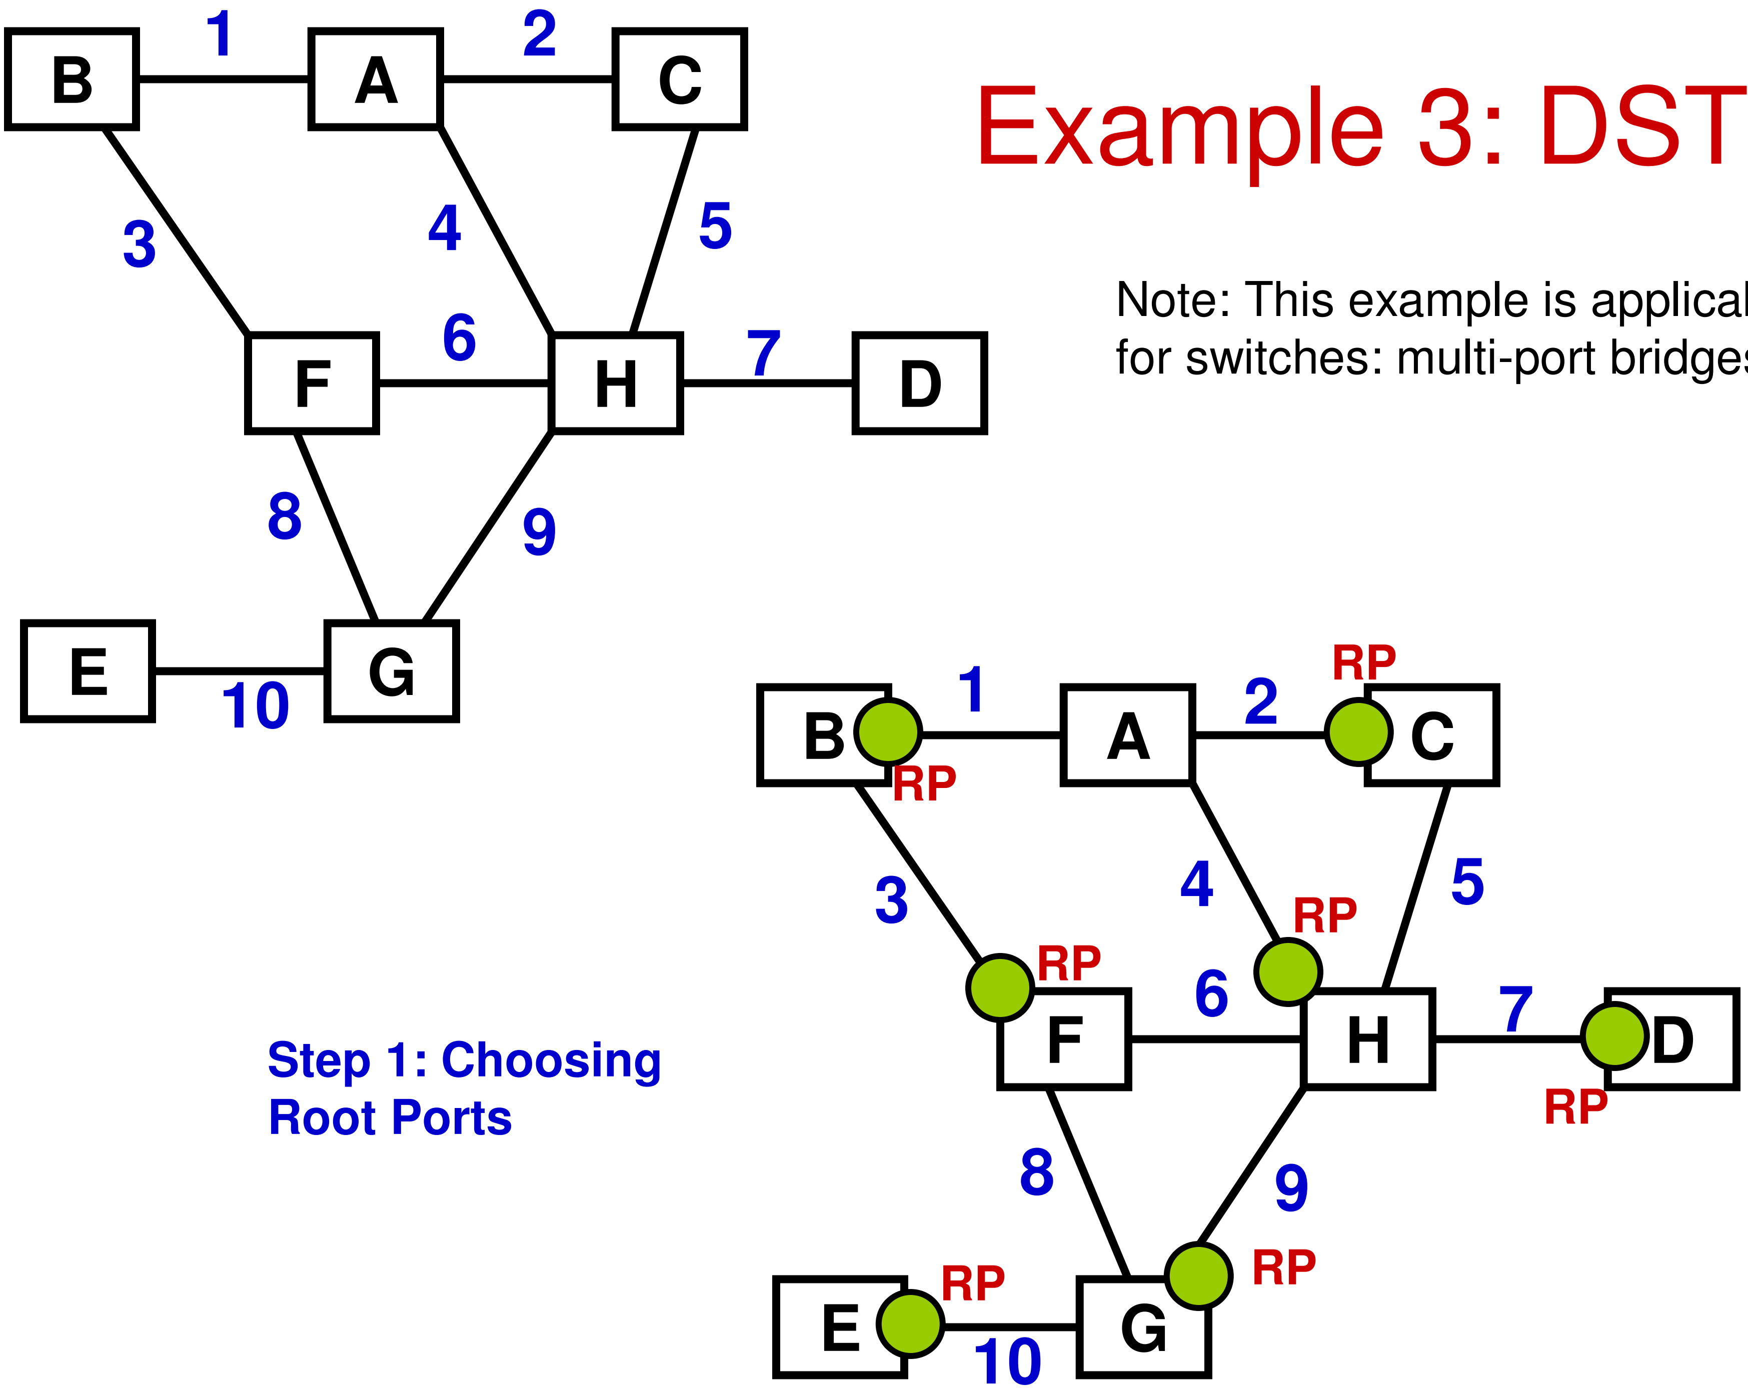

Step 1: Choosing
Root Ports

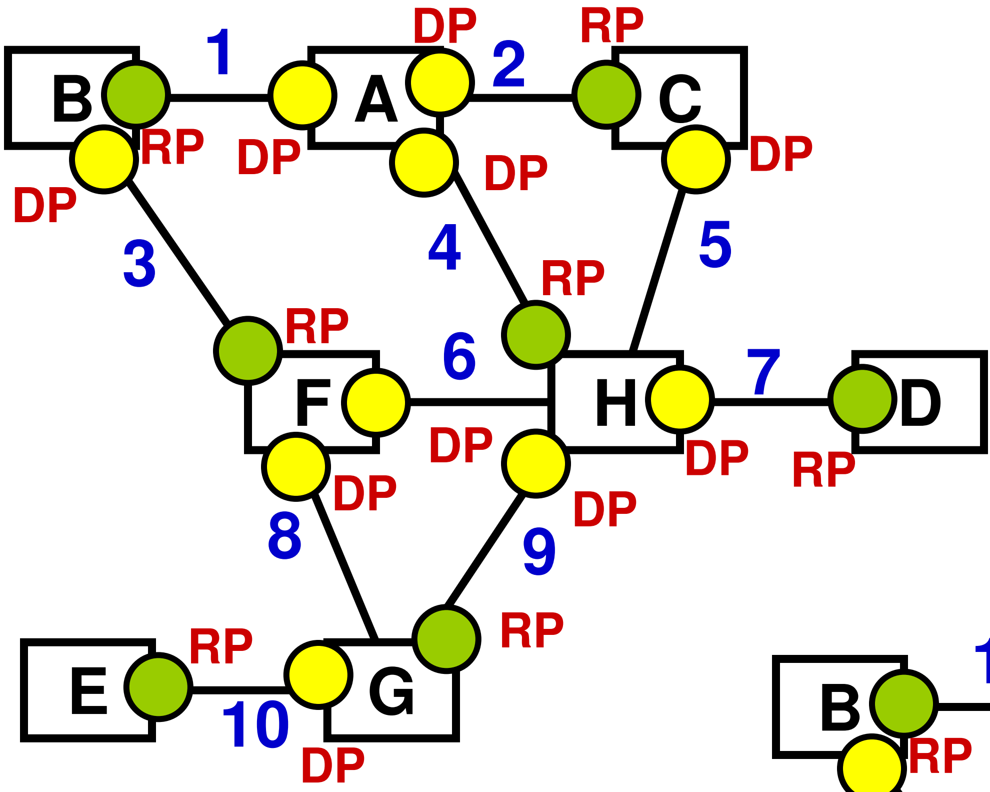

Step 2: Choosing Designated Ports

Forwarding Bridges for Link/LAN

1	A	6	F
2	A	7	H
3	B	8	F
4	A	9	H
5	C	10	G

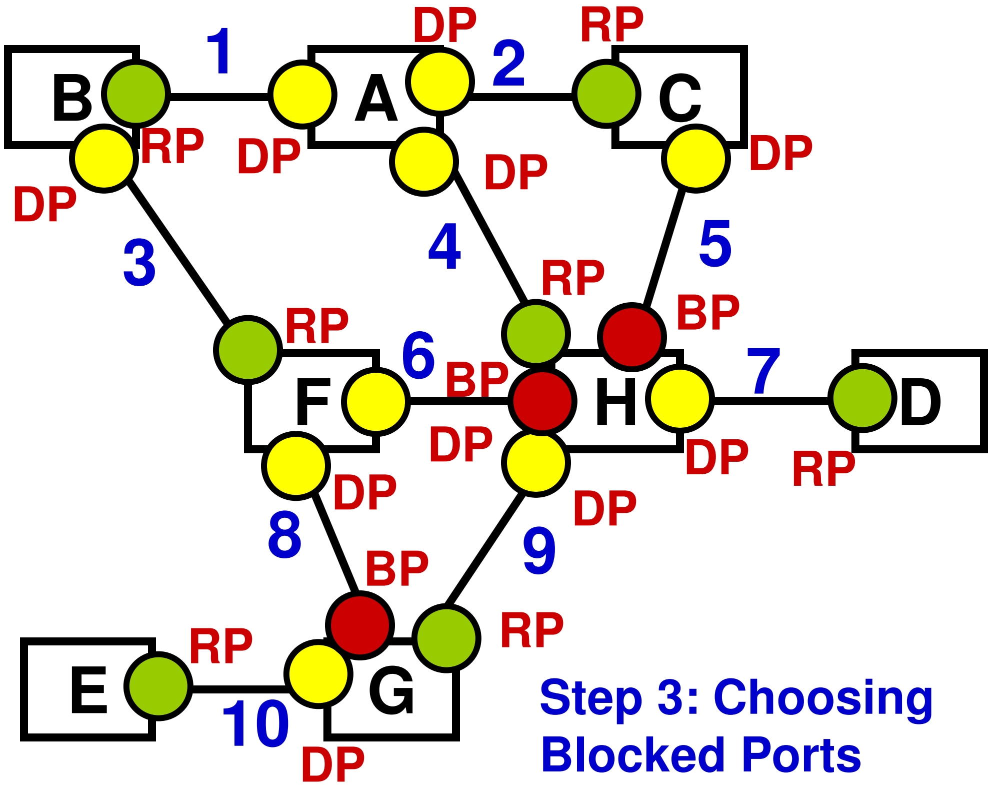

Step 3: Choosing Blocked Ports

Final DST

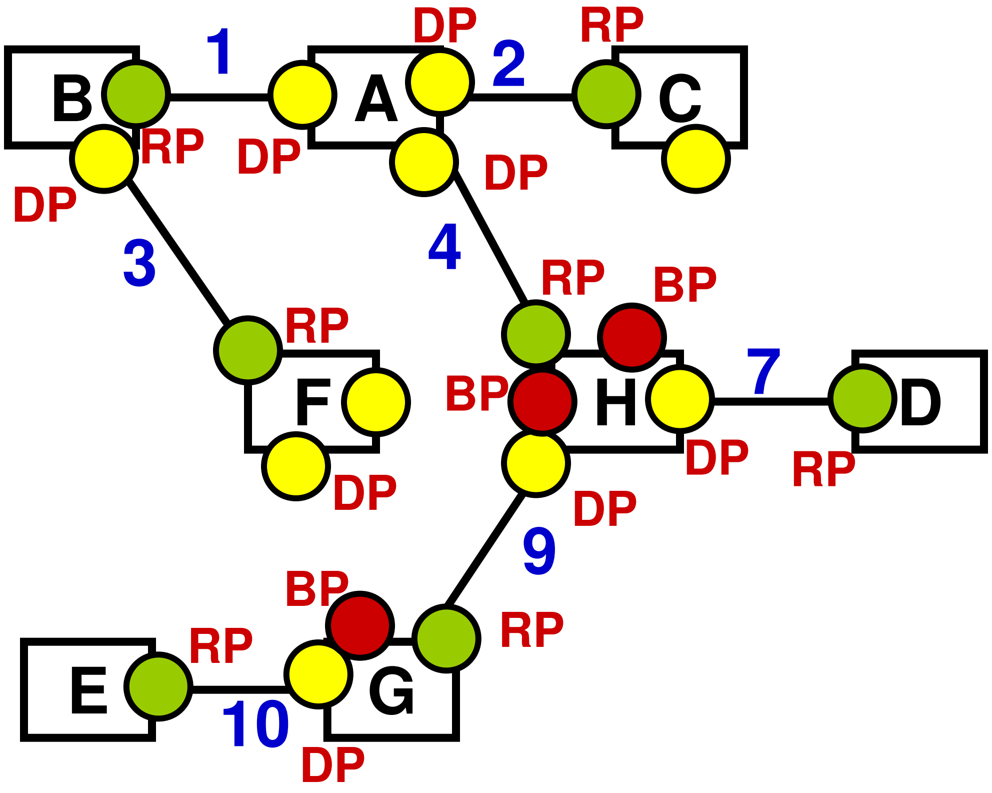
B
1
A
DP
2
RP
C
RP
DP
DP
DP
3
4
RP
BP
RP
F
BP
H
7
D
DP
DP
RP
DP
9
BP
RP
E
RP
G
10
DP

Switch

- A switch is a multi-port bridge
- If R is the rate at which a computer can send data, and N is the number of computers connected to a switch, then the maximum throughput possible is RN/2.
- On the other hand, only a pair of computers can communicate in a hub system at any time. Hence the throughput is R.

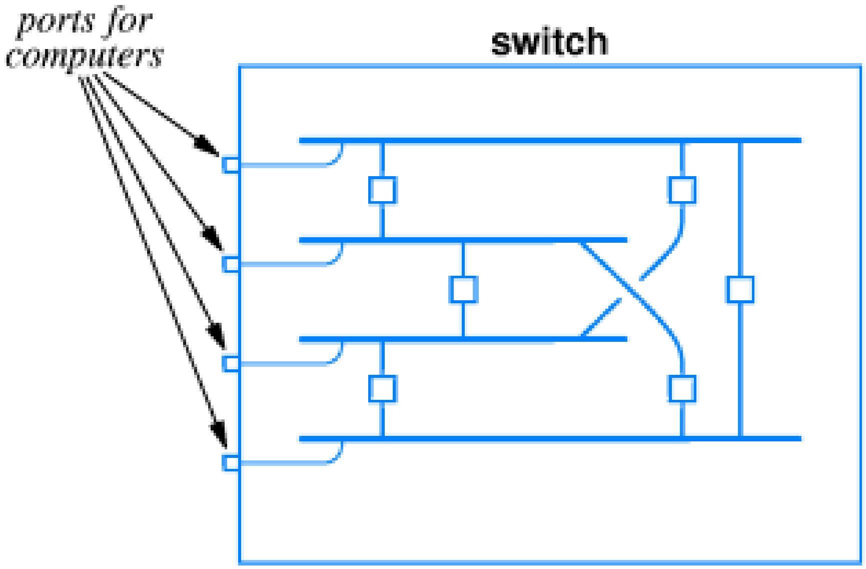

To emulate a N-port switch, we would need {N*(N-1)/2} bridges

Example 1: Number of bridges to emulate a 5-port switch is 5(5-1)/2 = 10.

Example 2: If 6 bridges are put together to emulate a switch, the number of ports needed for the switch is:
$N(N-1)/2 = 6 \Rightarrow N(N-1) = 12$
$\Rightarrow N^2 - N - 12 = 0 \Rightarrow (N - 4)(N + 3) = 0$
$\Rightarrow$ N = 4 or N = - 3. As N ≥ 0, N = 4 ports

Combining Switches and Hubs

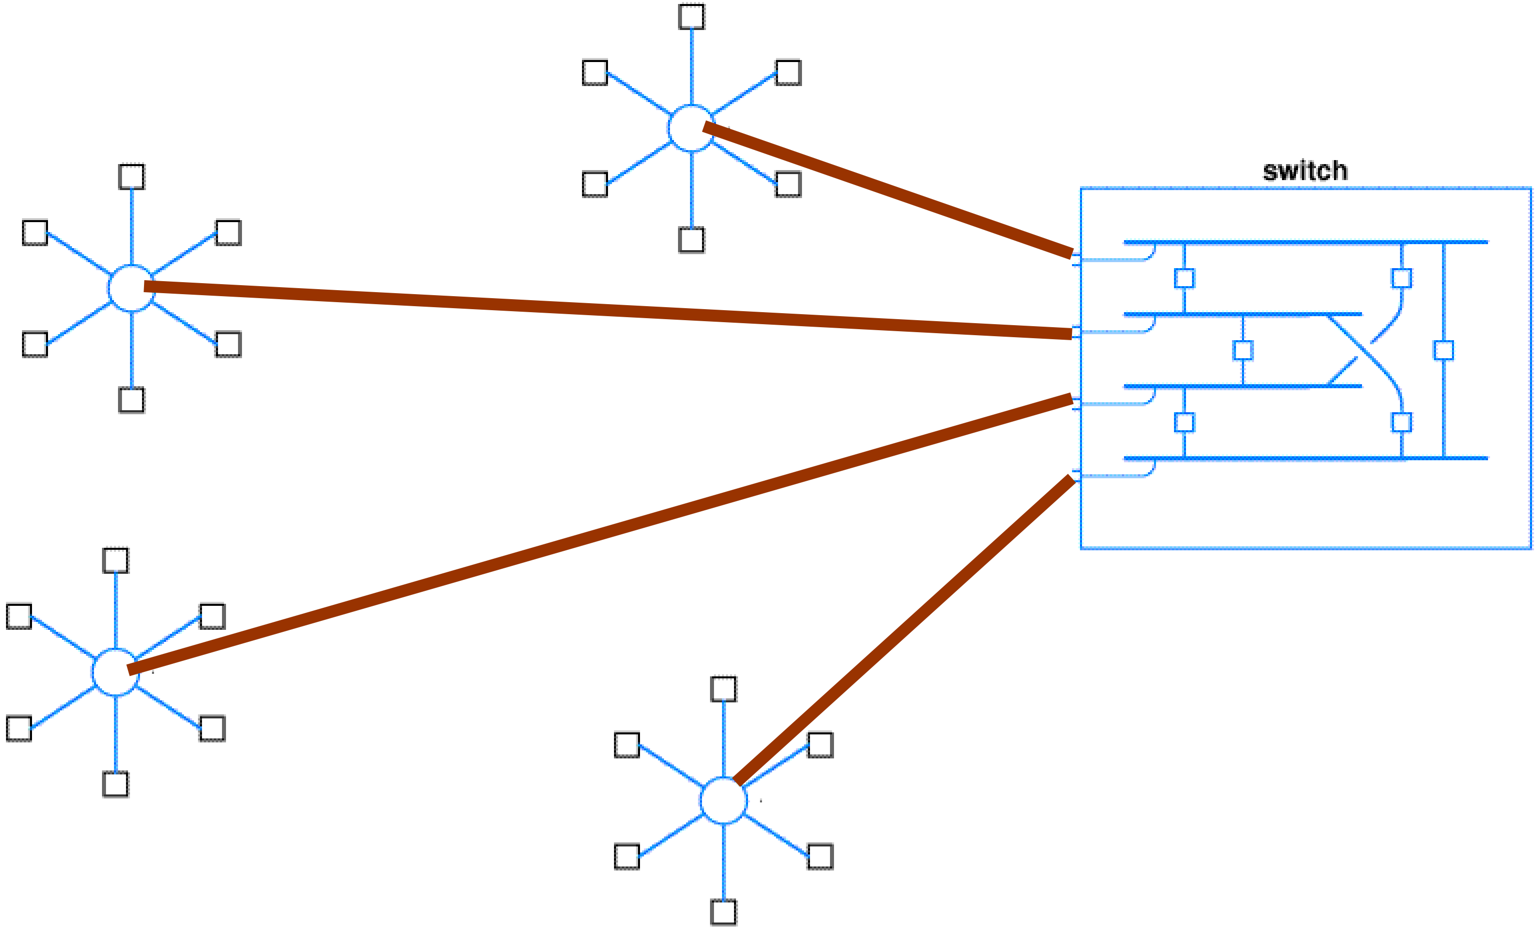

TCP/IP Stack Layers and the Devices

- **Repeater and Hub** operate at the Physical layer as all they do is to forward the signal received from one port to all the other ports.
- They do not filter any messages based on their destination address and neither they have an address for their own.
- For moderate and high traffic, a hub-or-repeater-based network will sustain a lower throughput. For low traffic, a hub-or-repeater-based network will perform better because there would be no table look-up delay and there is no significant channel access delay.
- **Bridge and Switch** operate at the Data Link layer as they are involved in transferring messages across the LAN segments of an organizational network.
- The devices can filter messages based on their destination MAC address and they themselves have a unique MAC address for each of their ports.
- For low traffic, the bridge/switch table lookup delay dominates over the channel access delay. However, for moderate and high traffic across an extended LAN, the filtering capability of these devices according to the destination MAC address helps to sustain a larger throughput compared to the hub-or-repeater-based networks.

4.3 VLANs

Virtual LANs (VLANs)

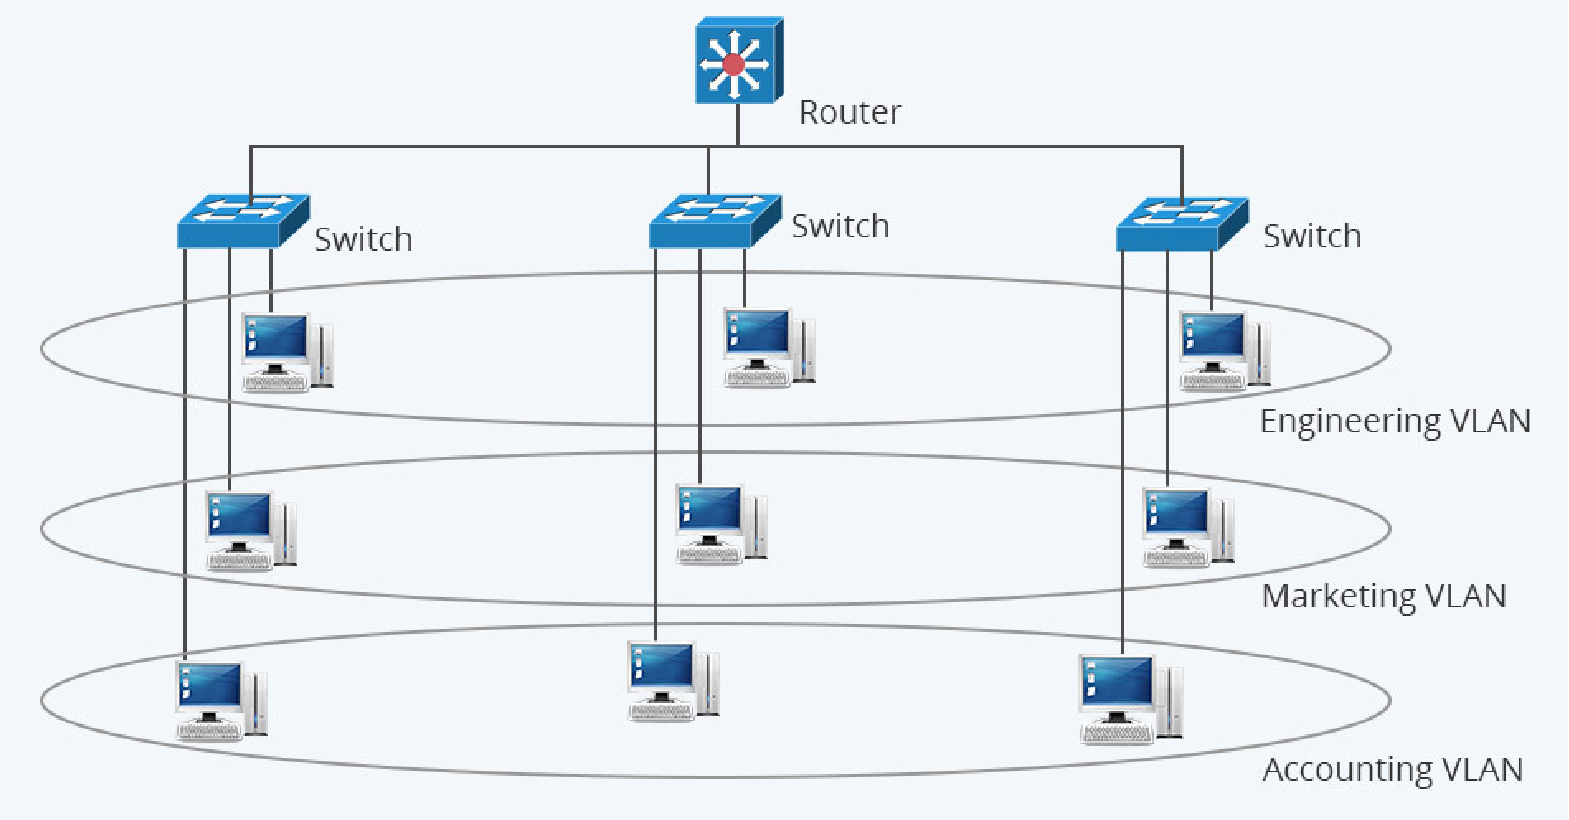

Source: http://www.fiber-optic-transceiver-module.com/vlan-vs-subnet.html

Virtual LANs (VLANs)

- VLANs are layer 2 constructs and they typically have a one-to-one relation with IP subnets that are layer 3 constructs.
- The idea is to partition a single LAN (broadcast domain) into multiple distinct broadcast domains, each of which is called a VLAN.
- We need a layer-3 device to forward packets from one VLAN to another VLAN.
- A layer-2 switch, by itself, cannot directly forward traffic from a computer in one VLAN to a computer in another VLAN. However, we could have the layer-2 switch interact with a router to facilitate this [Router in a Stick model].
- The IP addresses to the network interfaces in each VLAN are assigned in such a way that they are in the appropriate broadcast domain.
- Each VLAN has a unique VLAN ID.
- If a connection needs to carry the traffic of only one VLAN, then we set the connection to be in “Access” mode (identified by the particular VLAN ID). Example: Computer to Switch
- If a connection needs to carry the traffic of multiple VLANs, then we set the connection to be in “Trunk” mode (indicating the different VLAN IDs in its configuration). Example: Switch to Switch or Switch to Router

VLANs: Router in a Stick

- It is possible to setup a switch to recognize different broadcast domains by configuring it with the different VLAN IDs. However, a layer-2 switch can handle communication only between machines in a particular VLAN.

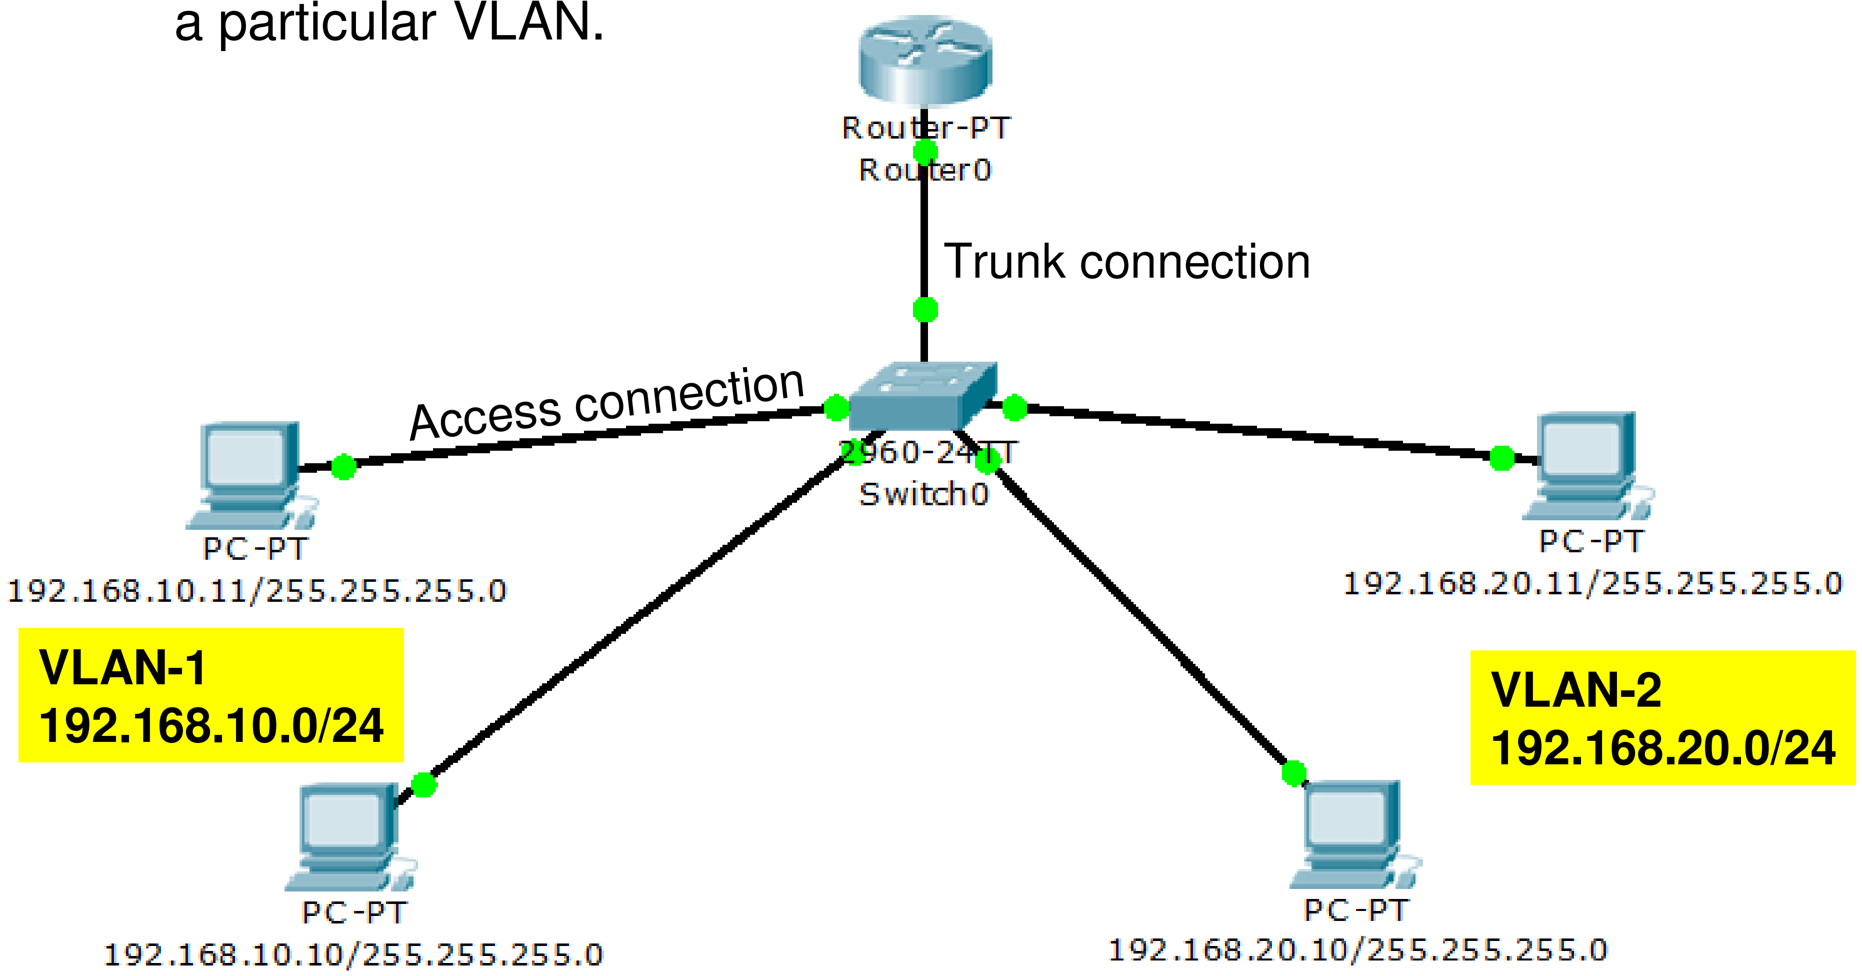

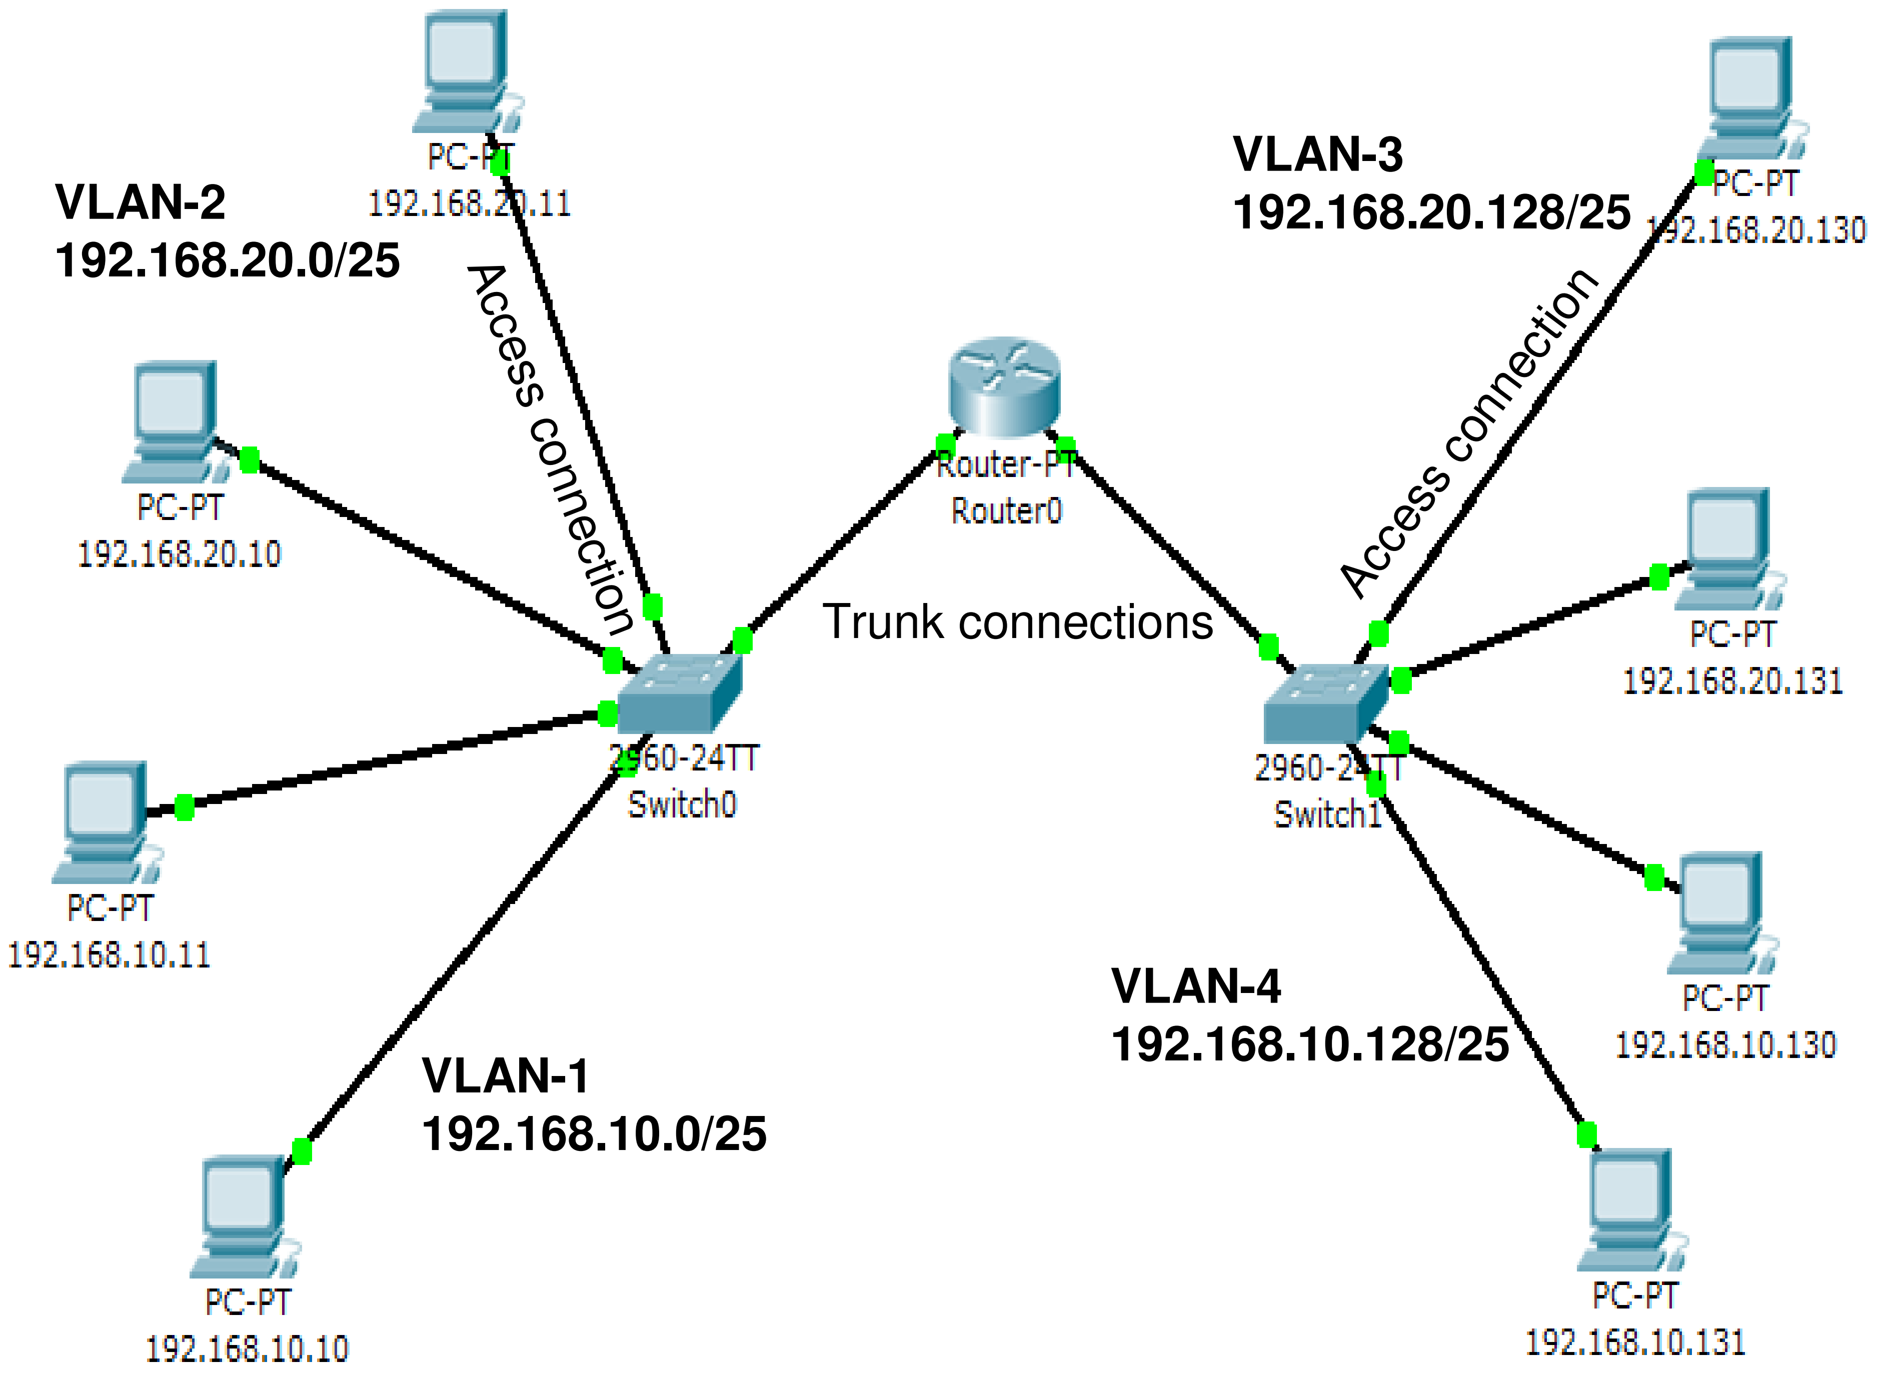
VLAN-2
192.168.20.0/25
PC-PT
192.168.20.11
Access connection
PC-PT
192.168.20.10
PC-PT
192.168.10.11
2960-24TT
Switch0
PC-PT
192.168.10.10
VLAN-1
192.168.10.0/25
Router-PT
Router0
Trunk connections
VLAN-3
192.168.20.128/25
PC-PT
192.168.20.130
Access connection
PC-PT
192.168.20.131
2960-24TT
Switch1
VLAN-4
192.168.10.128/25
PC-PT
192.168.10.130
PC-PT
192.168.10.131

VLANs Spread across Switches

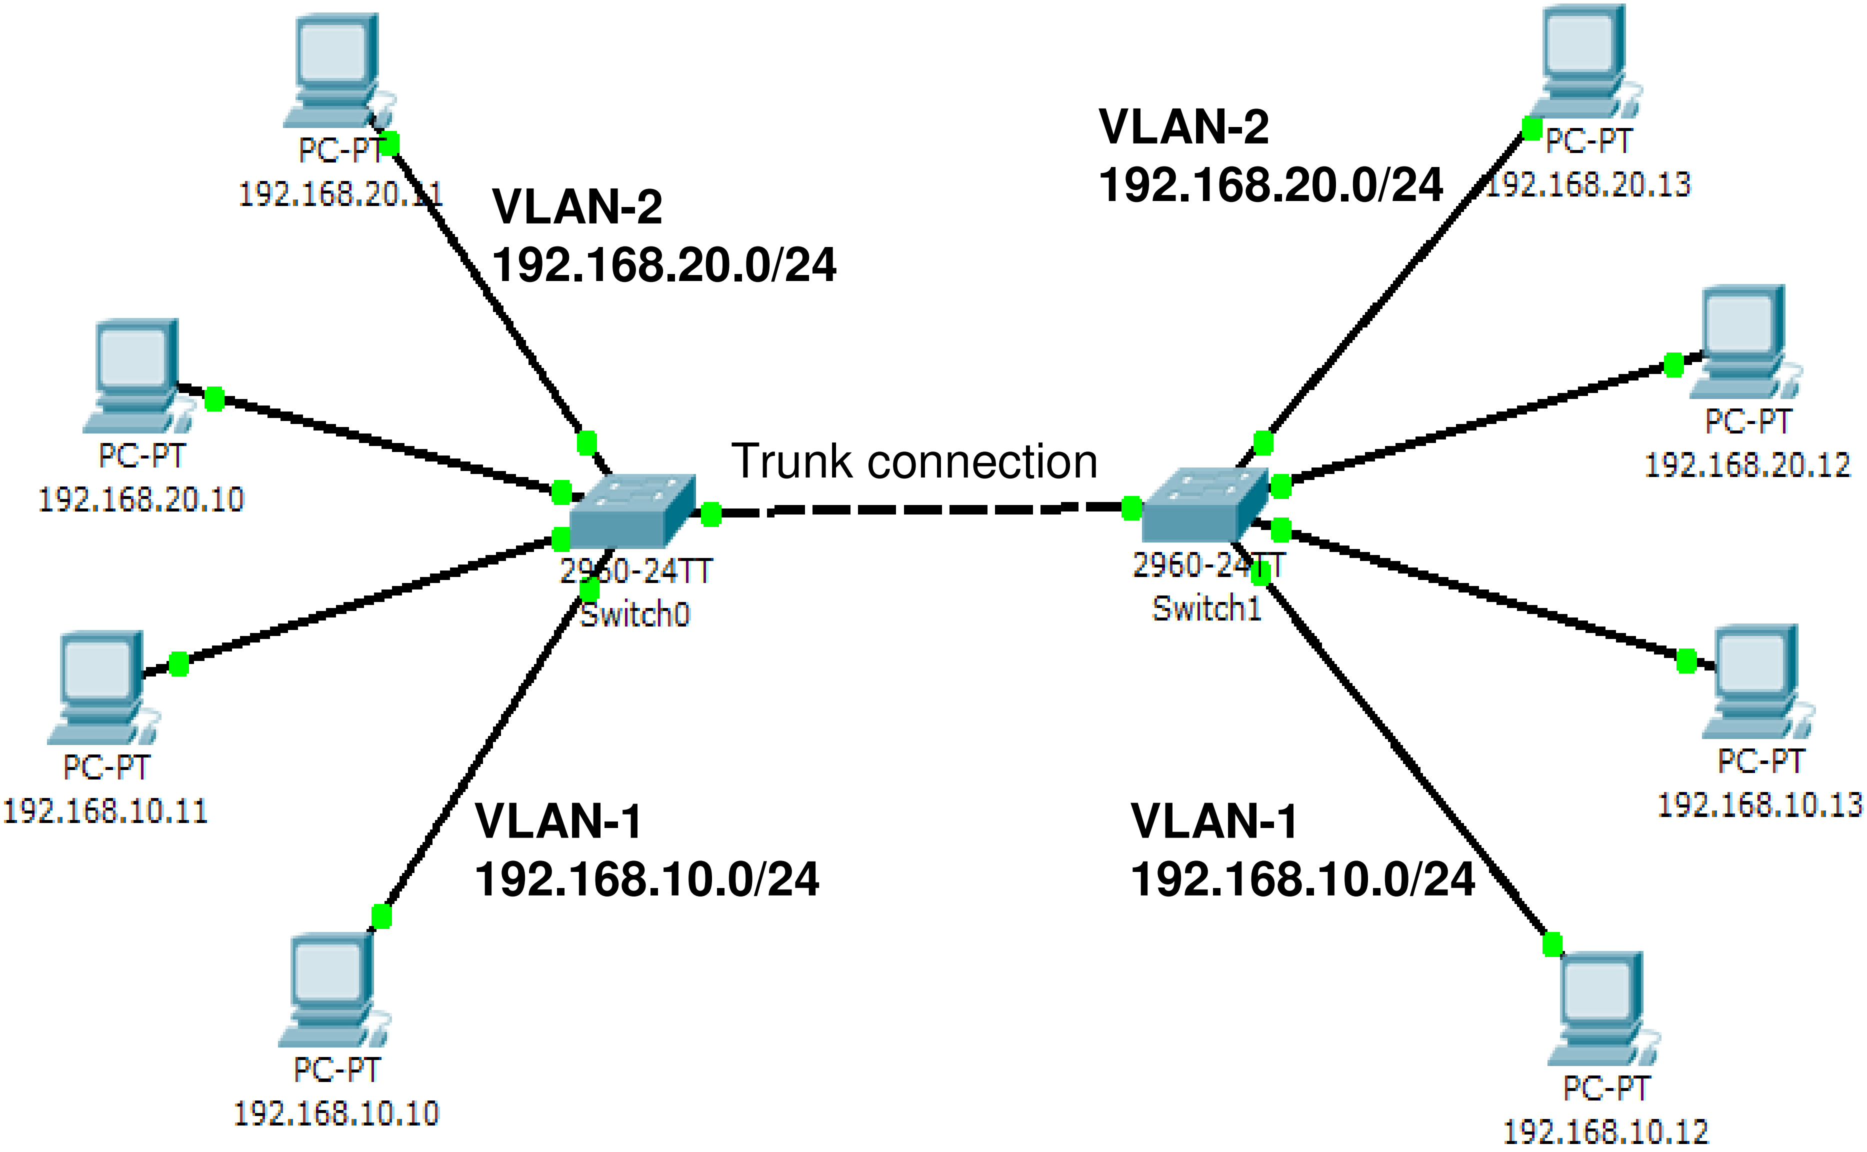

VLAN-2
192.168.20.0/24

VLAN-2
192.168.20.0/24

Trunk connection

VLAN-1
192.168.10.0/24

VLAN-1
192.168.10.0/24