CMPE 344 Computer Networks
Spring 2019

Internetworking
Part 1

Reading: Peterson and Davie, §3.1

26/10/2019 1

Not all networks are directly connected

- Limit to how many hosts can be attached
  - Point-to-point: Two hosts
  - Ethernet: 1,024 hosts
- Limit to how large of a geographic area a single network can serve
  - Ethernet: 2,500 m
  - Wireless: Limited by radio range
  - Point-to-point: Long, but do not serve the area
- Our major goal: Build global networks
  - Enable communication between hosts that are not directly connected

2

A switch provides a star topology

- Switched networks are scalable!
- It is possible for many hosts to transmit at the full link speed (bandwidth) provided that the switch is designed with enough aggregate capacity

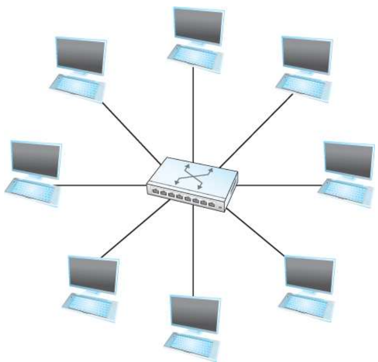

3

Input-output ports of a switch

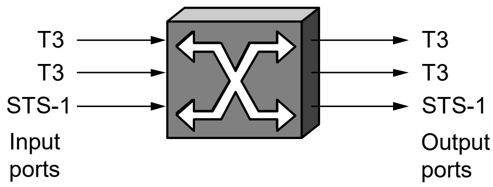

<u>Some typical leased line services:</u>
T1 (or DS1): 1.544 Mbps (24 digital voice circuits)
T3 (or DS3): 44.736 Mbps (28 T1 links)
STS-1 (or OC-1): 51.840 Mbps
STS-3 (or OC-3): 155.250 Mbps
STS: Synchronous Transport Signal
OC: Optical Carrier

4

Switching or forwarding

- Switching or forwarding is the main function of the network layer
- A switch's primary job is to receive incoming packets on one of its links and to transmit them on some other link
- A switch looks at the header of the packet (frame) for an identifier or an address
  – e.g., 48-bit Ethernet address
- Three approaches to switching:
  – Datagram or connectionless
  – Virtual circuit or connection-oriented
  – Source routing (less common)

5

Switch architecture

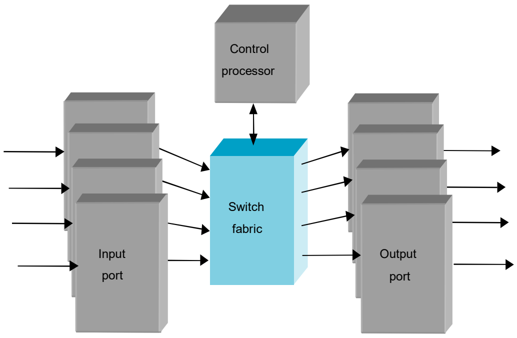

6

Types of switching fabrics

- Shared memory
  - Memory bandwidth determines the throughput
- Shared bus
  - Bus bandwidth determines the throughput
- Crossbar
  - Fast
- Self-routing
  - Fast and scalable

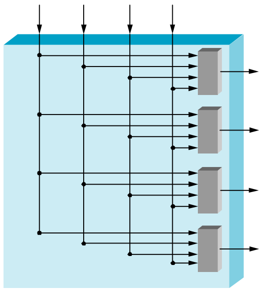

7

Connectionless (datagram) networks: Datagram switching

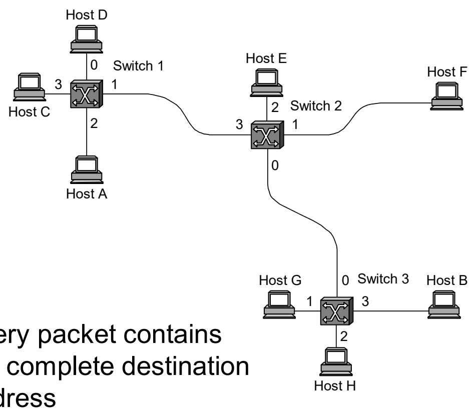

Every packet contains the complete destination address
Used in IP networks

Destination	Port
A	3
B	0
C	3
D	3
E	2
F	1
G	0
H	0

Forwarding (routing) table at switch 2

8

Datagram switching

- A routing algorithm builds the forwarding (routing) tables
- No connection state needs to be established before the first packet is sent
- No way of knowing if the packet can be successfully delivered
- Each packet is forwarded independent of previous packets that might have been sent to the same destination
- A switch or link failure might not have any serious effect on communication if it is possible to find an alternate route around the failure

9

Virtual circuit (VC) networks: VC switching

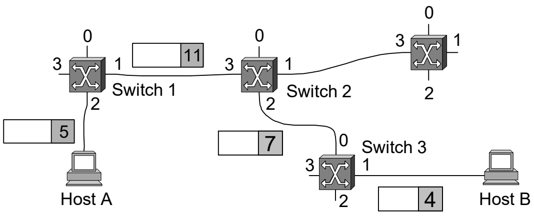

In IF	In VCI	Out IF	Out VCI
2	5	1	11

VC table entry at switch 1

In IF	In VCI	Out IF	Out VCI
3	11	2	7

VC table entry at switch 2

In IF	In VCI	Out IF	Out VCI
0	7	1	4

VC table entry at switch 3

Connection-oriented approach: Requires that a VC from source to destination is set up before any data is sent.
Used in ATM, Frame Relay networks
Note that connection-oriented does NOT imply reliable!

10

More on virtual circuits

- Virtual circuit identifiers (VCIs) have link-local scope
  - Incoming and outgoing VCIs are not necessarily the same
- Whenever a new connection is created, a new VCI must be assigned to it on each link it will traverse; the assigned VCI value must not be in use
- Types of VCs
  - Permanent VCs (PVC): Set up by the administrator
  - Switched VCs (SVC): Dynamically set up by signaling
- If a switch or link in a connection fails, the circuit is broken!
- It is possible to allocate resources when VC is set up
  - If there are not enough resources, connection request can be rejected

11

Source routing

- Information about network topology that is required to switch a packet across the network is provided by the source host in packet header

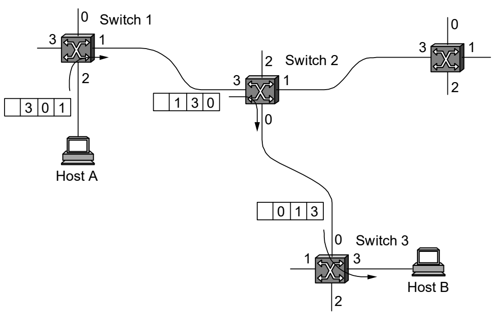

12

LAN switching

- LAN switches or bridges
  - A bridge is a switch: multi-input multi-output device
  - A single Ethernet segment: 10 Mbps
  - Ethernet bridge with n ports: Up to 10n Mbps
  - Frames with destinations on the same segment need not be forwarded onto other ports

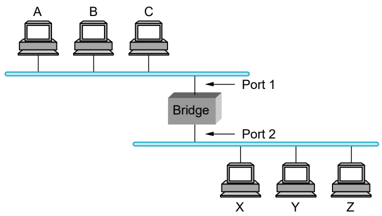

An extended LAN with
One broadcast domain
Two collision domains

13

Learning bridges

- Bridges can build forwarding tables themselves
  - Inspect the source (MAC) address in the frames received and record the port received
  - Each entry has an associated timeout

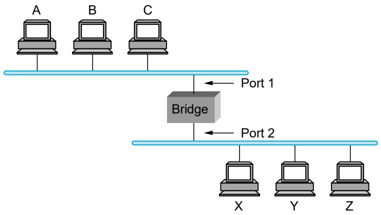

Host	Port
A	1
B	1
C	1
X	2
Y	2
Z	2

14

Loops in the topology

- Frame F with unknown destination: Use flooding
- On seeing F_2, Bridge 1 copies it onto LAN 1
- Similarly, Bridge 2 copies F_1 to LAN 1
- Now, these copies are forwarded to LAN 2. This cycle goes forever

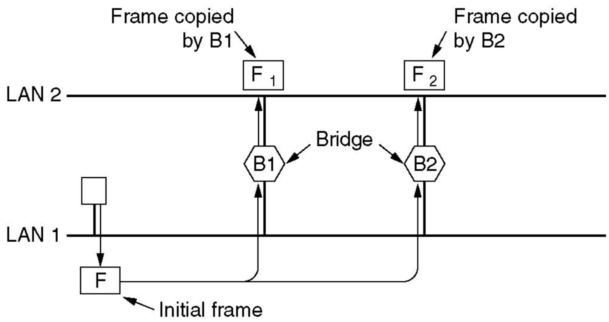

15

Spanning tree bridges

- Extended LAN with loops and the corresponding spanning tree

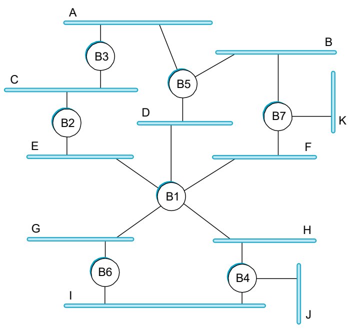

Frames can loop in LAN forever!

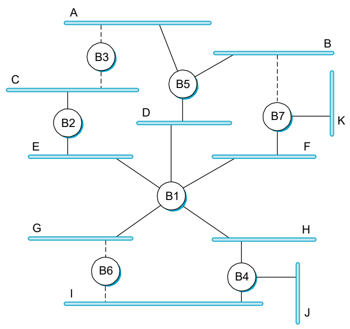

No cycles

16

Spanning tree algorithm

- Bridges choose one bridge as the root
  - This choice is made after each bridge broadcasts its (unique) serial number
  - The bridge with the lowest serial number becomes the root
- Tree of shortest paths from the root to every bridge and LAN is constructed
- If a bridge or LAN fails, a new tree is constructed (algorithm continues to run)
- The distributed algorithm is standardized in IEEE 802.1D

17

"Translating" bridges

- Devices that translate between two LAN technologies: e.g., 802.x to 802.y

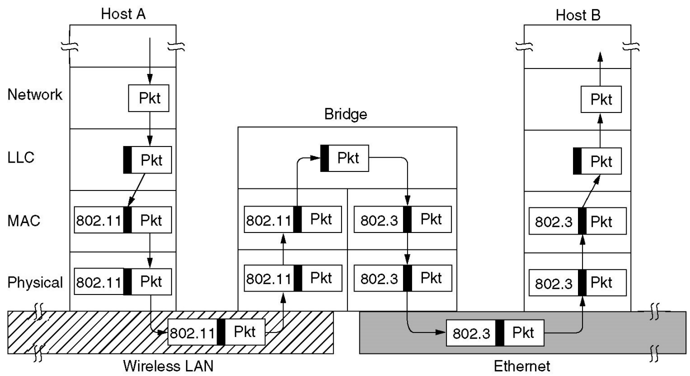

18

Virtual LANs

- Single extended LAN partitioned into several seemingly separate LANs
- Each VLAN has an id (or color)

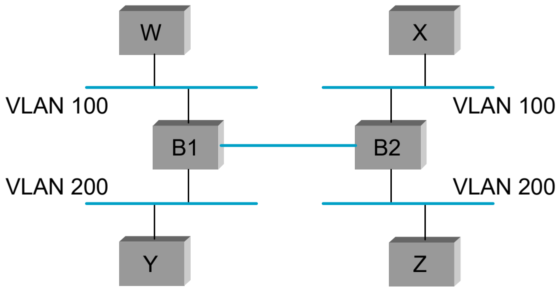

19

More on VLANs

- In absence of VLANs, any broadcast packet will reach all hosts (one broadcast domain)
- Let's suppose W and X are in VLAN 100; Y and Z in VLAN 200
  - VLAN IDs on each port of bridges B1 and B2 are configured (Link B1-B2 in both VLANs)

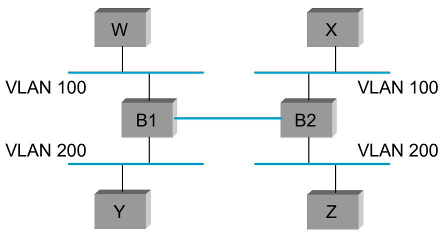

A broadcast packet sent by X will be forwarded by B2 to B1, but not to Z. B1 will forward to W, but not to Y

20

Cell switching and ATM networks

- ATM (Asynchronous transfer mode)
  – Connection-oriented, packet-switched technology
  – Uses virtual circuits
  – Signaling or connection setup phase: Resources are allocated at the switches along the circuit to ensure a particular QoS
  – Fixed-length packets called cells (48 bytes payload + 5 bytes header = 53 bytes)

ATM cell format

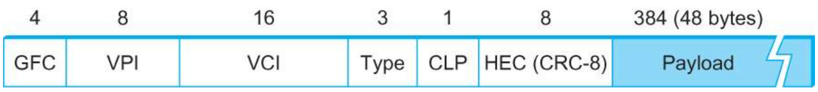

– User-Network Interface (UNI)
  - Host-to-switch format
  - GFC: Generic Flow Control
  - VPI: Virtual Path Identifier
  - VCI: Virtual Circuit Identifier
  - Type: management, congestion control
  - CLP: Cell Loss Priority
  - HEC: Header Error Check (CRC-8)

– Network-Network Interface (NNI)
  - Switch-to-switch format
  - GFC becomes part of VPI field

ATM cells

- Advantages of cells
  - facilitate the implementation of hardware switches
  - enable parallelism
- Disadvantage of cells
  - Header overhead: 48/53 = 90.6% efficient
- Performance advantage:
  - Queues of cells tend to be a little shorter
  - Shorter queues mean less delay for all traffic

23

Virtual paths

- 24-bit identifier split into two parts
  - 8-bit virtual path identifier (VPI)
  - 16-bit virtual circuit identifier (VCI)

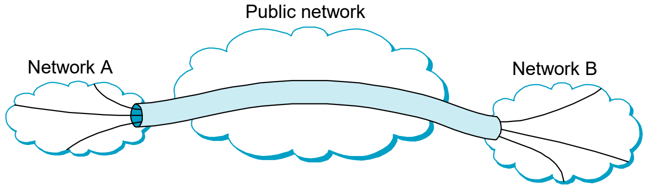

Much less connection-state information stored in the switches, avoiding the need for big, expensive tables of per VCI information in the public network

24

Comparison of datagram and VC networks

Issue	Datagram nets	VC nets
Connection setup	Not needed	Required
Addressing	Each packet contains full source-destination addresses	Each packet contains a short VCI
State information	Switches do not keep connection state info	Switches must keep connection state of each VC
Forwarding	Each packet is forwarded independently	All packets follow the route chosen when VC is set up
Effect of switch failures	None (except for a few packet losses at point of failure)	All VCs passing through must be terminated
QoS	Difficult	Easy (if enough resources can be allocated in advance)

25