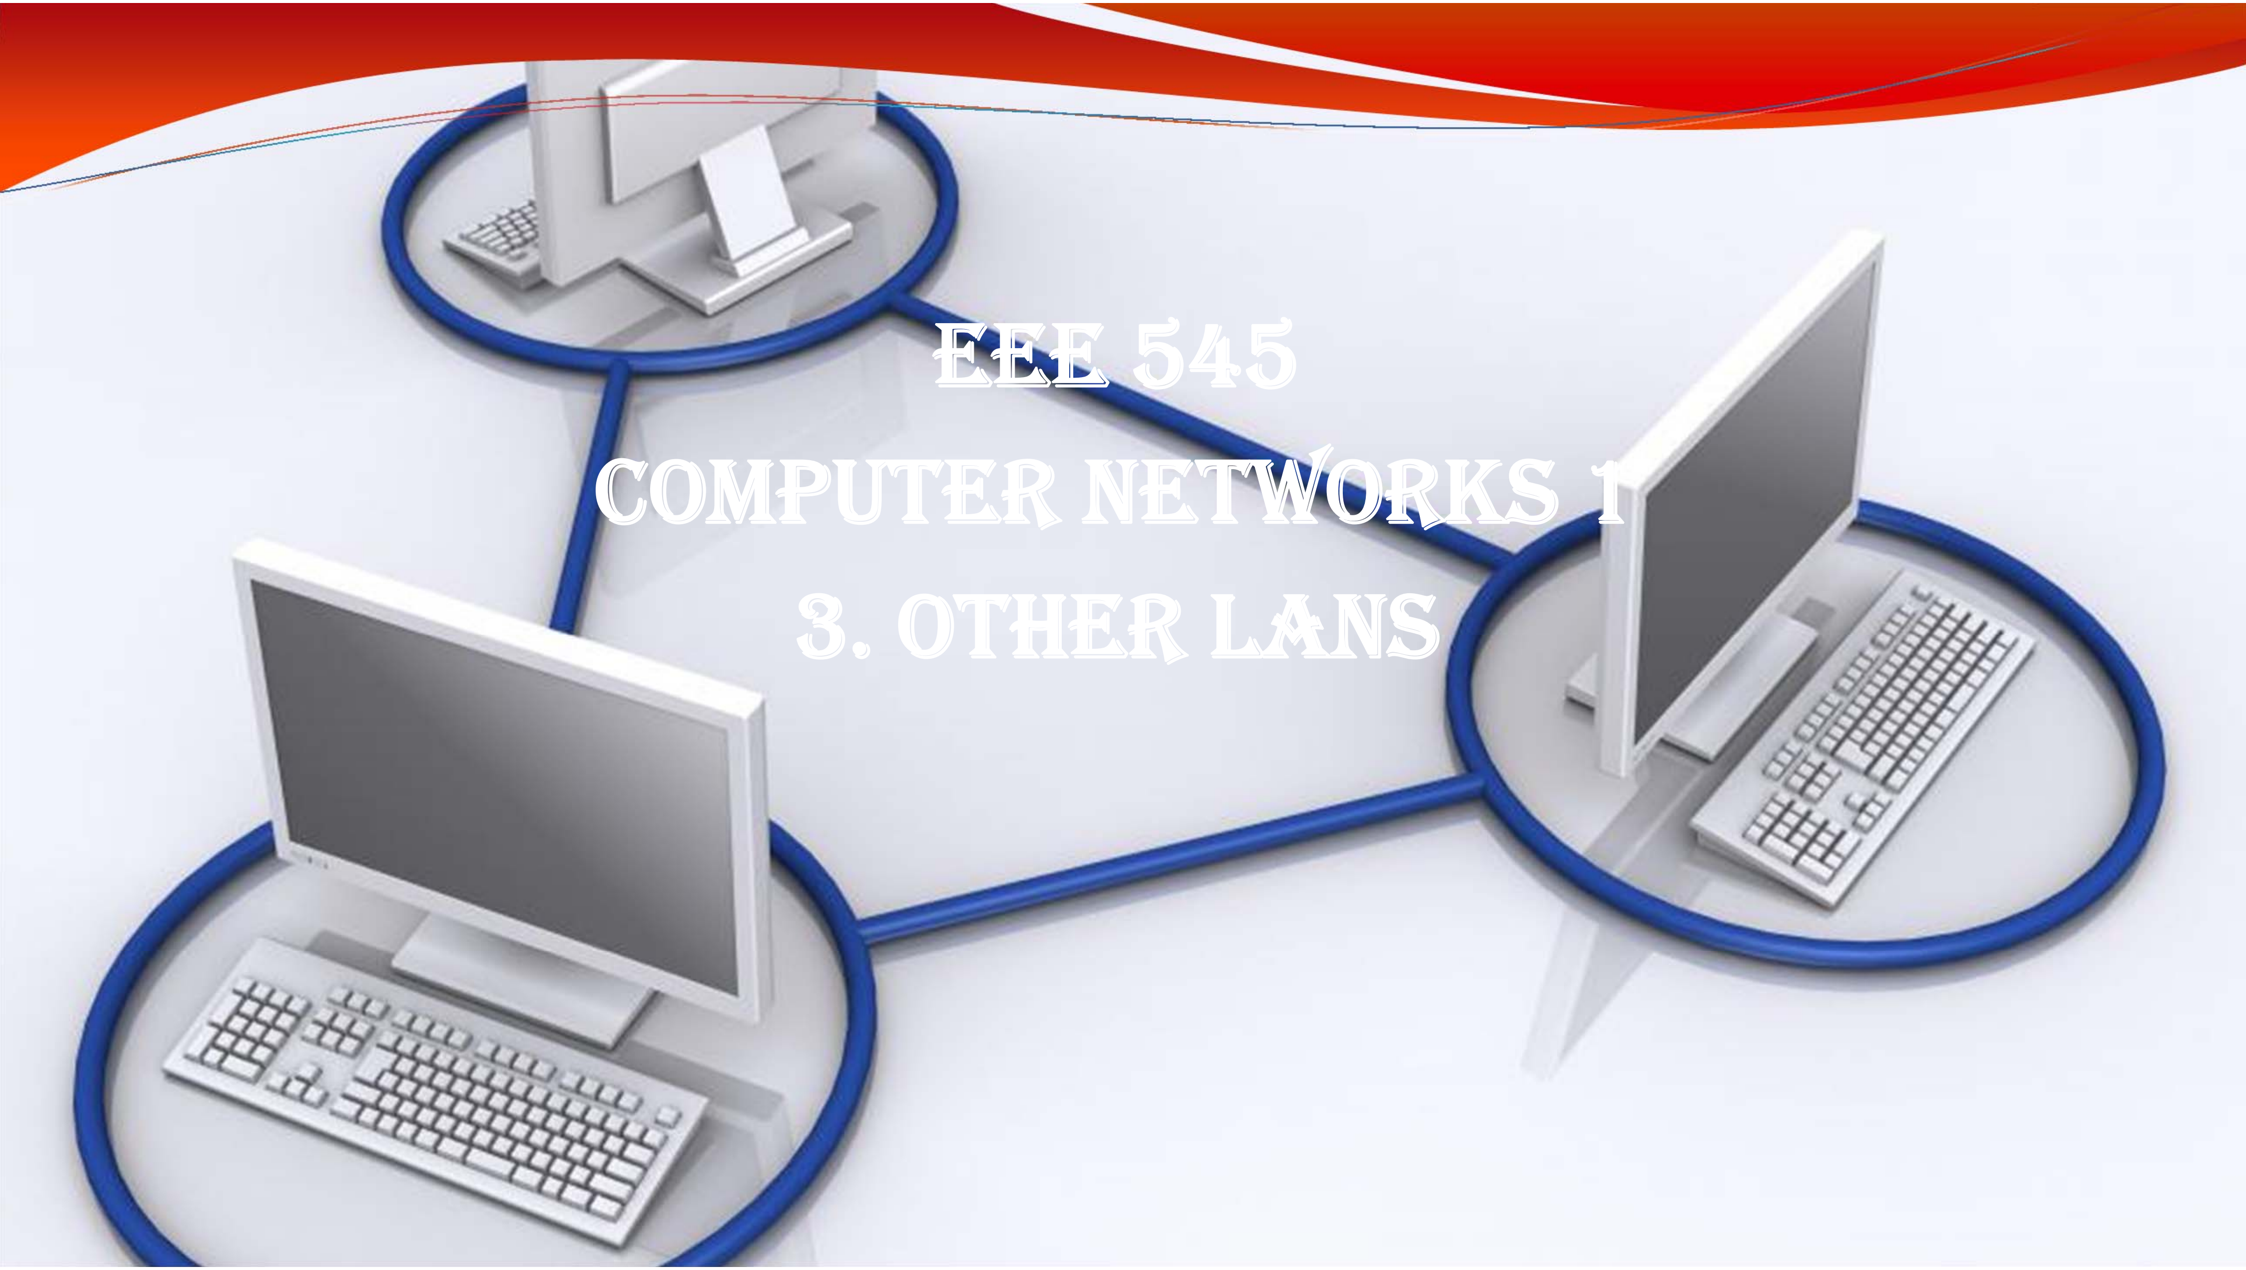

EEE 545

Computer Networks 1

3. Other LANs

3.1 Virtual Private Network (VPN)

- Suppose you want to connect to your workplace network from home. Your workplace, however, has a security policy that does not allow "outside" IP addresses to access essential internal resources. How do you proceed, without leasing a dedicated telecommunications line to your workplace?
- A virtual private network, or VPN, provides a solution; it supports creation of virtual links that join far-flung nodes via the Internet.
- Your home computer creates an ordinary Internet connection (TCP or UDP) to a workplace VPN server.
- Each end of the connection is associated with a software-created virtual network interface; each of the two virtual interfaces is assigned an IP address. When a packet is to be sent along the virtual link, it is actually encapsulated and sent along the original Internet connection to the VPN server, wending its way through the commodity Internet; this process is called tunnelling. To all intents and purposes, the virtual link behaves like any other physical link.

3.1 Virtual Private Network (VPN)

- At the workplace side, the virtual network interface in the VPN server is attached to a router or switch; at the home user's end, the virtual network interface can now be assigned an internal workplace IP address.
- The home computer is now, for all intents and purposes, part of the internal workplace network. In the diagram shown, the user's regular Internet connection is via hardware interface eth0. A connection is established to Site A's VPN server; a virtual interface tun0 is created on the user's machine which appears to be a direct link to the VPN server.
- The tun0 interface is assigned a Site-A IP address. Packets sent via the tun0 interface in fact travel over the original connection via eth0 and the Internet.

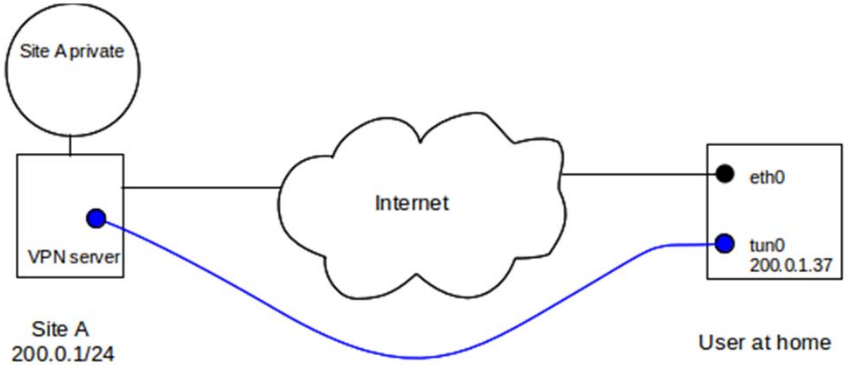

VPN: blue link represents *tunnel*. Actual connection is made via eth0
The tun0 interface is a virtual network interface with a Site-A address

3.2 Carrier Ethernet

- Carrier Ethernet is a leased-line point-to-point link between two sites, where the subscriber interface at each end of the line looks like Ethernet (in some flavor).
- The physical path in between sites, however, need not have anything to do with Ethernet; it may be implemented however the carrier wishes. In particular, it will be (or at least appear to be) full-duplex, it will be collision-free, and its length may far exceed the maximum permitted by any IEEE Ethernet standard.

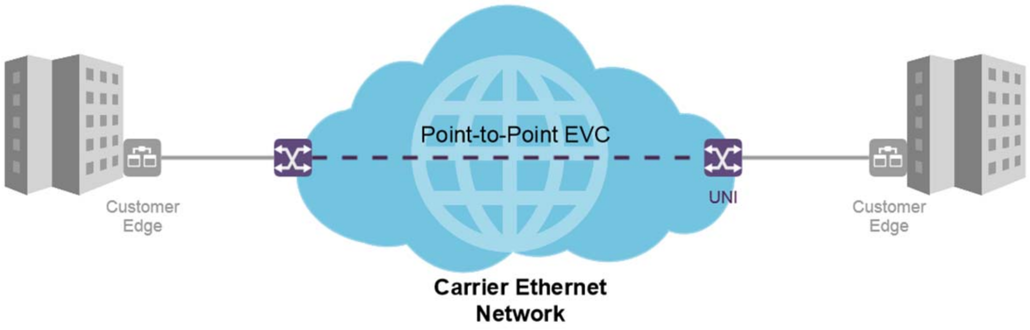

Carrier Ethernet

3.3 Token Ring

- It may come as a surprise that there is a simple multiple-access mechanism that is not only collision-free, but which supports fairness in the sense that if N stations wish to send then each will receive 1/N of the opportunities.
- That method is Token Ring. Actual implementations come in several forms, from Fiber-Distributed Data Interface (FDDI) to so-called “IBM Token Ring”. The central idea is that stations are connected in a ring.
- Packets will be transmitted in one direction. Stations in effect forward most packets around the ring, although they can also remove a packet. (It is perhaps more accurate to think of the forwarding as representing the default cable connectivity; non-forwarding represents the station’s momentarily breaking that connectivity.)

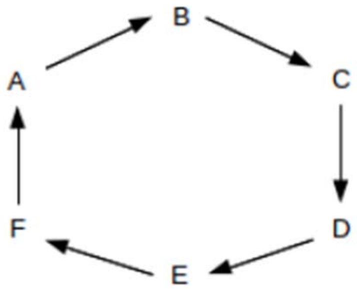

Token Ring

3.3 Token Ring

- When the network is idle, all stations agree to forward a special, small packet known as a token. When a station, say A, wishes to transmit, it must first wait for the token to arrive at A. Instead of forwarding the token, A then transmits its own packet; this travels around the network and is then removed by A.
- At that point (or in some cases at the point when A finishes transmitting its data packet) A then forwards the token. If all stations have packets to send, then we will have something like the following:
- • A waits for the token
- • A sends a packet
- • A sends the token to B
- • B sends a packet
- • B sends the token to C
- • C sends a packet
- • C sends the token to D
- • …

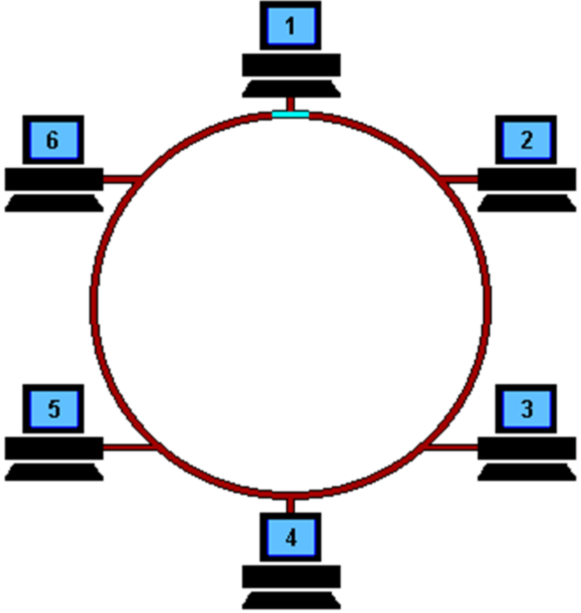

+ All stations get an equal number of chances to transmit, and no bandwidth is wasted on collisions.

- When stations are powered off it is essential that the packets continue forwarding
- Some station has to watch out in case the token disappears, or in case a duplicate token appears.

3.4 Virtual Circuits

- Virtual-circuit switching (or routing) is an alternative to datagram switching. **In datagram switching, routers know the next_hop to each destination, and packets are addressed by destination. In virtual-circuit switching, routers know about end-to-end connections, and packets are "addressed" by a connection ID.**
- Before any packets can be sent, a connection needs to be established first. For that connection, the route is computed and then each link along the path is assigned a connection ID, traditionally called the VCI, for Virtual Circuit Identifier. In most cases, VCIs are only locally unique; that is, the same connection may use a different VCI on each link. The lack of global uniqueness makes VCI allocation much simpler. ,
- Although the VCI keeps changing along a path, the VCI can still be thought of as identifying the connection. To send a packet, the host marks the packet with the VCI assigned to the host–router1 link.
- Packets arrive at (and depart from) switches via one of several ports, which we will assume are numbered beginning at 0. Switches maintain a connection table indexed by <VCI, port> pairs; unlike a forwarding table, the connection table has a record of every connection through that switch at that particular moment.
- As a packet arrives, its inbound VCIin and inbound port in are looked up in this table; this yields an outbound <VCIout, portout> pair. The VCI field of the packet is then rewritten to VCIout, and the packet is sent via portout.

3.4 Virtual Circuits

- As an example, consider the network below. Switch ports are numbered 0,1,2,3. Two paths are drawn in, one from A to F in red and one from B to D in green; each link is labelled with its VCI number in the same colour.

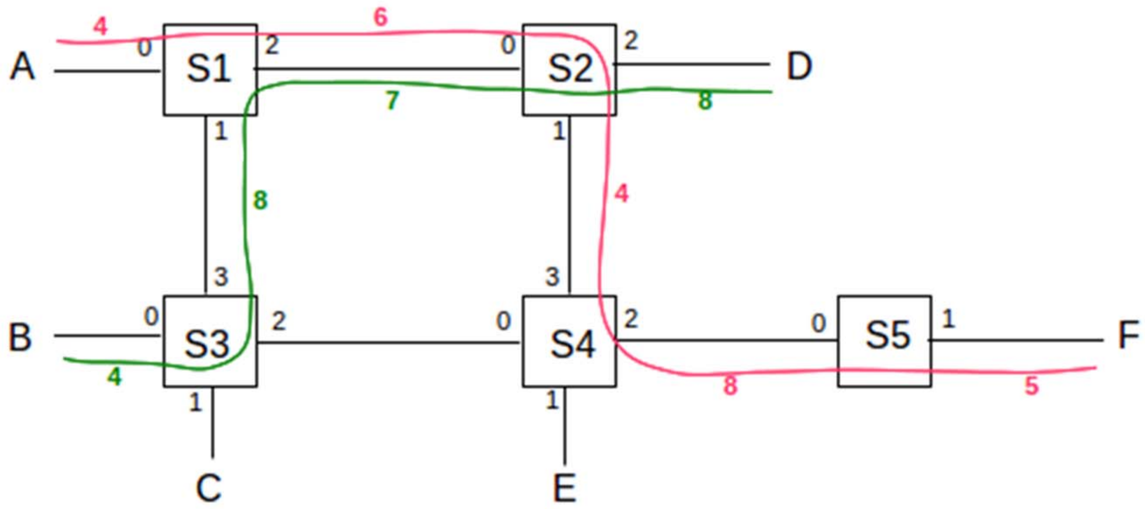

3.4 Virtual Circuits

- We will construct virtual-circuit connections between
- • A and F (shown above in red)
- • A and E
- • A and C
- • B and D (shown above in green)
- • A and F again (a separate connection)
- The following VCIs have been chosen for these connections. The choices are made more or less randomly here, **but in accordance with the requirement that they be unique to each link.** Because links are generally taken to be bidirectional, a VCI used from S1 to S3 cannot be reused from S3 to S1 until the first connection closes.

3.4 Virtual Circuits

- • A to F: A-4-S1-6-S2-4-S4-8-S5-5-F; this path goes from S1 to S4 via S2
- • A to E: A-5-S1-6-S3-3-S4-8-E; this path goes, for no particular reason, from S1 to S4 via S3, the opposite corner of the square
- • A to C: A-6-S1-7-S3-3-C
- • B to D: B-4-S3-8-S1-7-S2-8-D
- • A to F: A-7-S1-8-S2-5-S4-9-S5-2-F (again, a separate connection)

!! It can be verified that no two different paths use the same VCI on any link.

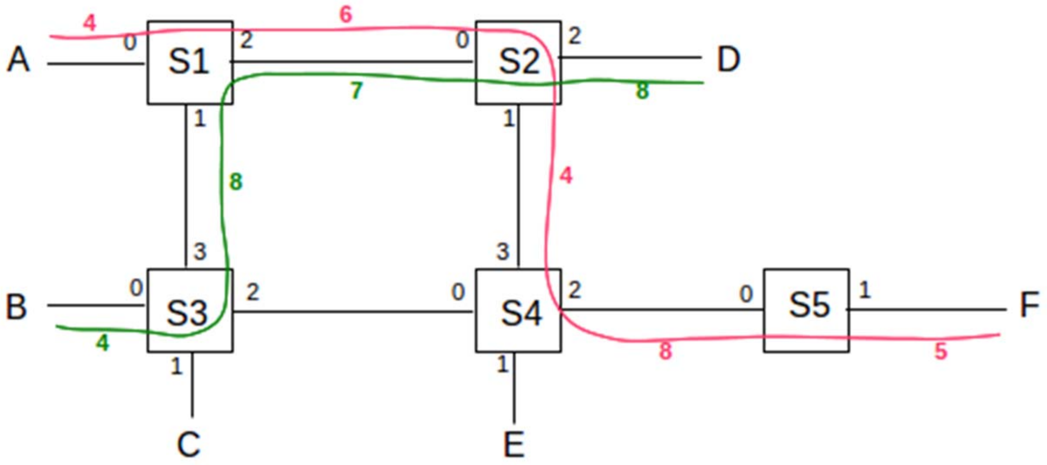

3.4 Virtual Circuits

- We now construct the actual <VCI,port> tables for the switches S1-S4

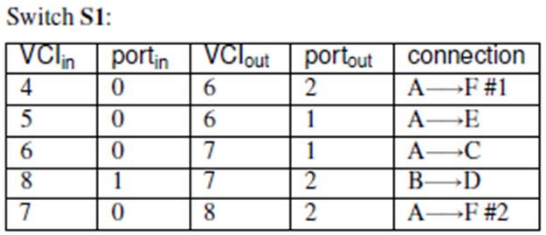

Switch **S1**:

VCI_{in}	$port_{in}$	VCI_{out}	$port_{out}$	connection
4	0	6	2	A ⟶ F #1
5	0	6	1	A ⟶ E
6	0	7	1	A ⟶ C
8	1	7	2	B ⟶ D
7	0	8	2	A ⟶ F #2

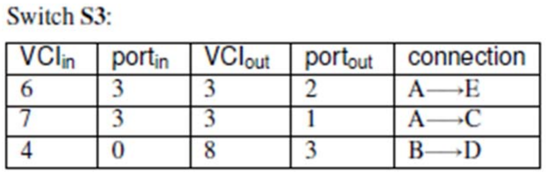

Switch **S3**:

VCI_{in}	$port_{in}$	VCI_{out}	$port_{out}$	connection
6	3	3	2	A ⟶ E
7	3	3	1	A ⟶ C
4	0	8	3	B ⟶ D

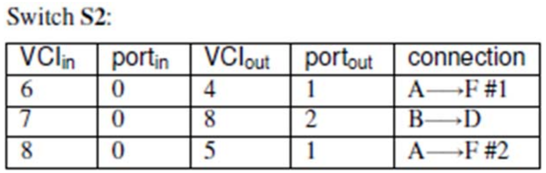

Switch **S2**:

VCI_{in}	$port_{in}$	VCI_{out}	$port_{out}$	connection
6	0	4	1	A ⟶ F #1
7	0	8	2	B ⟶ D
8	0	5	1	A ⟶ F #2

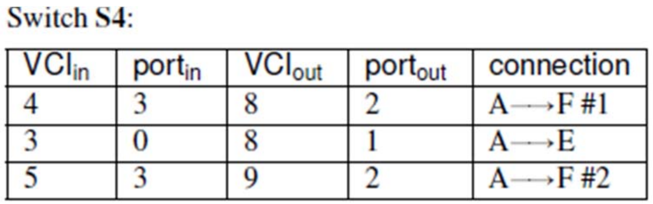

Switch **S4**:

VCI_{in}	$port_{in}$	VCI_{out}	$port_{out}$	connection
4	3	8	2	A ⟶ F #1
3	0	8	1	A ⟶ E
5	3	9	2	A ⟶ F #2

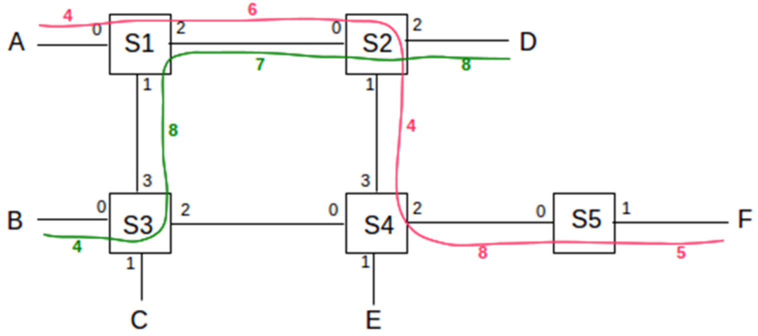

3.4 Virtual Circuits

Virtual-circuit switching offers the following advantages:

- Connections can get quality-of-service guarantees, because the switches are aware of connections and can reserve capacity at the time the connection is made
- Headers are smaller, allowing faster throughput
- Headers are small enough to allow efficient support for the very small packet sizes that are optimal for voice connections. ATM packets, for instance, have 48 bytes of data.

Datagram forwarding, on the other hand, offers these advantages:

- Routers have less state information to manage.
- Router crashes and partial connection state loss are not a problem.
- If a router or link is disabled, rerouting is easy and does not affect any connection state.
- Per-connection billing is very difficult.

3.6 Adventures in Radioland

3.6.1 Collisions

- One fundamental change is that Ethernet-like collision detection is no longer feasible over radio. This has to do with the relative signal strength of the remote signal at the local transmitter.
- Wi-Fi also supports its PCF mode that involves fewer (but not zero) collisions through the use of central-point polling. Finally, WiMAX and LTE switch from polling to scheduling to further reduce collisions, though the potential for collisions is still inevitable when new stations join the network.
- It is also worth pointing out that, while an Ethernet collision affects every station in the physical Ethernet (the “collision domain”), wireless collisions are often **local**.

3.6.2 Hidden Nodes

- In wireless communication, two nodes A and B that are not in range of one another – and thus cannot detect one another – may still have their signals interfere at a third node C. This creates an additional complication to collision handling.

3.6.3 Bandwith

- To radio engineers, “band width” means the frequency range used by a signal, not the data transmission rate.
- No information can be conveyed using a single frequency; even signalling by switching a carrier frequency off and on at a low rate “blurs” the carrier into a band of nonzero width.
- In keeping with this we will for the remainder of this chapter use the term “data rate” for what we have previously called “bandwidth”. We will use the terms “channel width” or “width of the frequency band” for the frequency range.

3.6.4 Cost

- Higher data rates require wider frequency bands. To reduce costs in the face of fixed demand, the usual strategy is to make the coverage zones smaller, either by reducing power (and adding more access points as appropriate), or by using directional antennas, or both.

3.6.5 Multipath

- While a radio signal generally covers a wide area – even with ordinary directional antennas – it does so in surprisingly non-uniform ways. A signal may reach a receiver through a line-of-sight path and also several reflected paths, possibly of varying length. In addition to reflection, the signal may be subject to reflection like scattering and diffraction. All of this together is known as multipath interference (or, if analog audio is involved, multipath distortion; in the analog TV era this was ghosting).

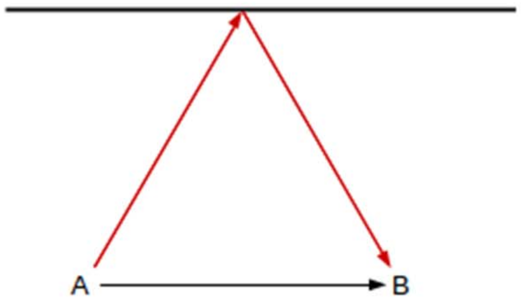

Line-of-sight and reflected signals

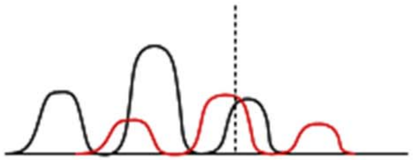

Superposition of encoded data