High Speed LANs

High Speed LANs

- Range of technologies
  - —Fast and Gigabit Ethernet
  - —Fibre Channel
  - —High Speed Wireless LANs

Why High Speed LANs?

- Office LANs used to provide basic connectivity
  - Connecting PCs and terminals to mainframes and midrange systems that ran corporate applications
  - Providing workgroup connectivity at departmental level
  - Traffic patterns were light
    - Emphasis on file transfer and electronic mail
- Speed and power of PCs has risen
  - Graphics-intensive applications and GUIs
- MIS organizations recognize LANs as essential
  - Began with client/server computing
    - Now dominant architecture in business environment
    - Intranetworks
    - Frequent transfer of large volumes of data

Applications Requiring High Speed LANs

- Centralized server farms
  - User needs to draw huge amounts of data from multiple centralized servers
  - E.g. Color publishing
    - Servers contain tens of gigabytes of image data
    - Downloaded to imaging workstations
- Power workgroups
- Small number of cooperating users
  - Draw massive data files across network
  - E.g. Software development group testing new software version or computer-aided design (CAD) running simulations
- High-speed local backbone
  - Processing demand grows
  - LANs proliferate at site
  - High-speed interconnection is necessary

Ethernet (CSMA/CD)

- Carriers Sense Multiple Access with Collision Detection
- Developed by Xerox – original Ethernet
- IEEE 802.3

IEEE802.3
Medium Access Control

- Random Access
  - Stations access medium randomly
- Contention
  - Stations content for time on medium

ALOHA

- Packet Radio
- When station has frame, it sends
- Station listens (for max round trip time) plus small increment
- If ACK, fine. If not, retransmit
- If no ACK after repeated transmissions, give up
- Frame check sequence (as in HDLC)
- If frame OK and address matches receiver, send ACK
- Frame may be damaged by noise or by another station transmitting at the same time (collision)
- Any overlap of frames causes collision
- Max utilization 18%

Slotted ALOHA

- Time in uniform slots equal to frame transmission time
- Need central clock (or other sync mechanism)
- Transmission begins at slot boundary
- Frames either miss or overlap totally
- Max utilization 37%

CSMA

- Propagation time is much less than transmission time
- All stations know that a transmission has started almost immediately
- First listen for clear medium (carrier sense)
- If medium idle, transmit
- If two stations start at the same instant, collision

- Wait reasonable time (round trip plus ACK contention)
- If no ACK then retransmit
- Max utilization depends on propagation time (medium length) and frame length
  - Longer frame and shorter propagation gives better utilization

CSMA Persistence and Backoff

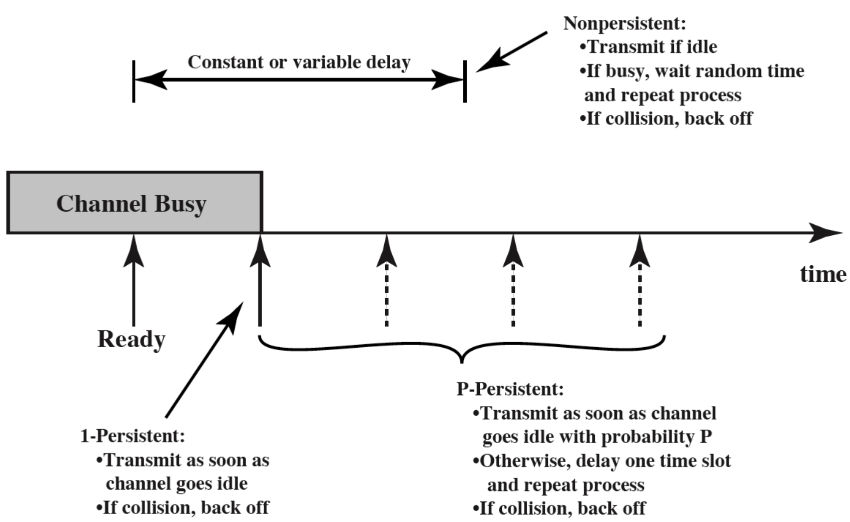

Nonpersistent CSMA

- Nonpersistent CSMA rules:
  1. if medium idle, transmit
  2. if medium busy, wait amount of time drawn from probability distribution (retransmission delay) & retry
- random delays reduces probability of collisions
- capacity is wasted because medium will remain idle following end of transmission
- nonpersistent stations are deferential

1-persistent CSMA

- To avoid idle channel time that non-persistent protocol used
- Station wishing to transmit listens and obeys following:
  1. If medium idle, transmit; otherwise, go to step 2
  2. If medium busy, listen until idle; then transmit immediately
- 1-persistent stations are selfish
- If two or more stations are waiting, collision guaranteed
  — Gets sorted out after collision

P-persistent CSMA

- Compromise that attempts to reduce collisions
  - Like nonpersistent
- And reduce idle time
  - Like1-persistent
- Rules:
  1. If medium idle, transmit with probability p, and delay one time unit with probability (1 – p)
     - Time unit typically maximum propagation delay
  2. If medium busy, listen until idle and repeat step 1
  3. If transmission is delayed one time unit, repeat step 1
- What is an effective value of p?

Value of p?

- Avoid instability under heavy load
- If *n* stations waiting to send then *np* stations will start transmitting
- If *np* > 1 on average there will be a collision
  - Repeated attempts to transmit almost guaranteeing more collisions
  - Retries compete with new transmissions
  - Eventually, all stations trying to send
    - Continuous collisions; zero throughput
- So *np* < 1 for expected peaks of *n*
  - If heavy load expected, *p* small
  - However, as *p* made smaller, stations wait longer
    - At low loads, this gives very long delays

CSMA/CD

- With CSMA, collision occupies medium for duration of transmission
- Stations listen whilst transmitting

1. If medium idle, transmit, otherwise, step 2
2. If busy, listen for idle, then transmit
3. If collision detected, jam then cease transmission
4. After jam, wait random time then start from step 1

CSMA/CD Operation

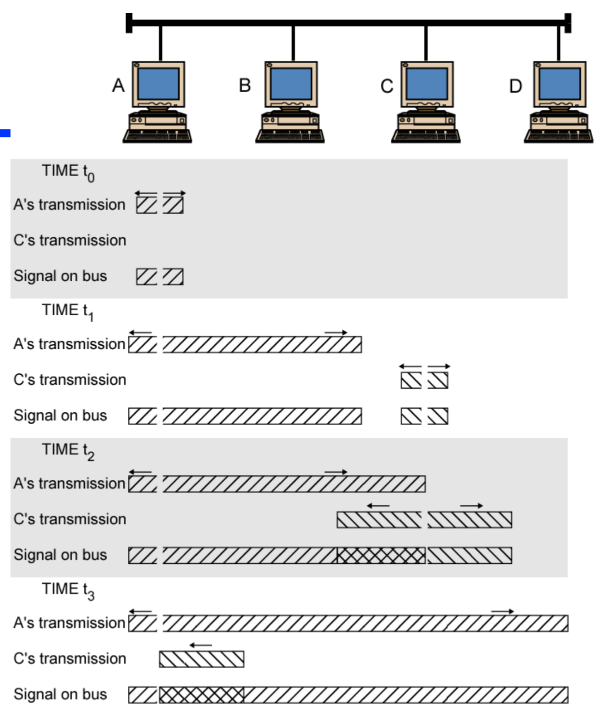

Which Persistence Algorithm?

- IEEE 802.3 uses 1-persistent
- Both nonpersistent and p-persistent have performance problems
- 1-persistent (p = 1) seems more unstable than p-persistent
  - —Greed of the stations
  - —But wasted time due to collisions is short (if frames long relative to propagation delay
  - —With random backoff, unlikely to collide on next tries
  - —**IEEE 802.3 and Ethernet use binary exponential backoff**

Binary Exponential Backoff

- Attempt to transmit repeatedly if repeated collisions
- First 10 attempts, mean value of random delay doubled
- Value then remains same for 6 further attempts
- After 16 unsuccessful attempts, station gives up and reports error
- As congestion increases, stations back off by larger amounts to reduce the probability of collision.
- 1-persistent algorithm with binary exponential backoff efficient over wide range of loads
  - —Low loads, 1-persistence guarantees station can seize channel once idle
  - —High loads, at least as stable as other techniques
- Backoff algorithm gives last-in, first-out effect
- Stations with few collisions transmit first

Collision Detection

- On baseband bus, collision produces much higher signal voltage than signal
- Collision detected if cable signal greater than single station signal
- Signal attenuated over distance
- Special collision presence signal

IEEE 802.3 Frame Format

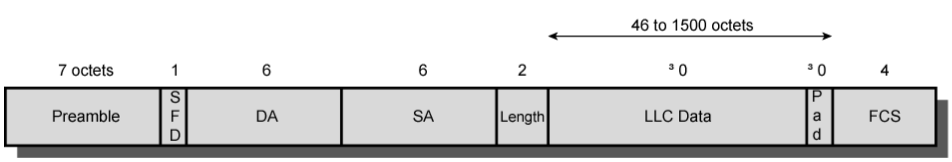

SFD = Start of frame delimiter
DA = Destination address
SA = Source address
FCS = Frame check sequence

100Mbps Fast Ethernet

- Use IEEE 802.3 MAC protocol and frame format
- 100BASE-X use physical medium specifications from FDDI
  - Two physical links between nodes
    - Transmission and reception
  - 100BASE-TX uses STP or Cat. 5 UTP
    - May require new cable
  - 100BASE-FX uses optical fiber
  - 100BASE-T4 can use Cat. 3, voice-grade UTP
    - Uses four twisted-pair lines between nodes
    - Data transmission uses three pairs in one direction at a time
- Star-wire topology
  - Similar to 10BASE-T

100Mbps (Fast Ethernet)

Table 16.3 IEEE 802.3 100BASE-T Physical Layer Medium Alternatives

	100BASE-TX		100BASE-FX	100BASE-T4
Transmission medium	2 pair, STP	2 pair, Category 5 UTP	2 optical fibers	4 pair, Category 3, 4, or 5 UTP
Signaling technique	MLT-3	MLT-3	4B5B, NRZI	8B6T, NRZ
Data rate	100 Mbps	100 Mbps	100 Mbps	100 Mbps
Maximum segment length	100 m	100 m	100 m	100 m
Network span	200 m	200 m	400 m	200 m

100BASE-X Data Rate and Encoding

- Unidirectional data rate 100 Mbps over single link
  - Single twisted pair, single optical fiber
- Encoding scheme same as FDDI
  - 4B/5B-NRZI
  - Modified for each option

100BASE-X Media

- Two physical medium specifications
- 100BASE-TX
  - Two pairs of twisted-pair cable
  - One pair for transmission and one for reception
  - STP and Category 5 UTP allowed
  - The MTL-3 signaling scheme is used
- 100BASE-FX
  - Two optical fiber cables
  - One for transmission and one for reception
  - Intensity modulation used to convert 4B/5B-NRZI code group stream into optical signals
  - 1 represented by pulse of light
  - 0 by either absence of pulse or very low intensity pulse

100BASE-T4

- 100-Mbps over lower-quality Cat 3 UTP
  - Taking advantage of large installed base
  - Cat 5 optional
  - Does not transmit continuous signal between packets
  - Useful in battery-powered applications
- Can not get 100 Mbps on single twisted pair
  - Data stream split into three separate streams
    - Each with an effective data rate of 33.33 Mbps
  - Four twisted pairs used
  - Data transmitted and received using three pairs
  - Two pairs configured for bidirectional transmission
- NRZ encoding not used
  - Would require signaling rate of 33 Mbps on each pair
  - Does not provide synchronization
  - Ternary signaling scheme (8B6T)

100BASE-T Options

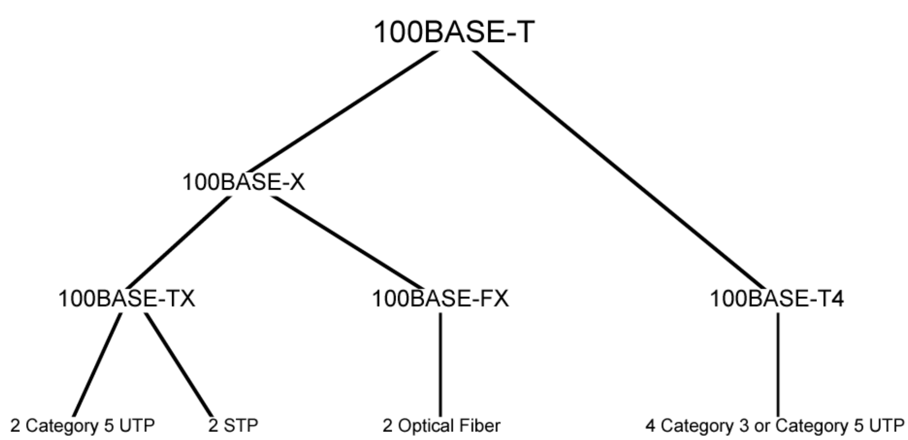

Full Duplex Operation

- Traditional Ethernet half duplex
  - Either transmit or receive but not both simultaneously
- With full-duplex, station can transmit and receive simultaneously
- 100-Mbps Ethernet in full-duplex mode, theoretical transfer rate 200 Mbps
- Attached stations must have full-duplex adapter cards
- Must use switching hub
  - Each station constitutes separate collision domain
  - In fact, no collisions
  - CSMA/CD algorithm no longer needed
  - 802.3 MAC frame format used
  - Attached stations can continue CSMA/CD

Mixed Configurations

- Fast Ethernet supports mixture of existing 10-Mbps LANs and newer 100-Mbps LANs
- E.g. 100-Mbps backbone LAN to support 10-Mbps hubs
  - Stations attach to 10-Mbps hubs using 10BASE-T
  - Hubs connected to switching hubs using 100BASE-T
    - Support 10-Mbps and 100-Mbps
  - High-capacity workstations and servers attach directly to 10/100 switches
  - Switches connected to 100-Mbps hubs using 100-Mbps links
  - 100-Mbps hubs provide building backbone
    - Connected to router providing connection to WAN

Gigabit Ethernet Configuration

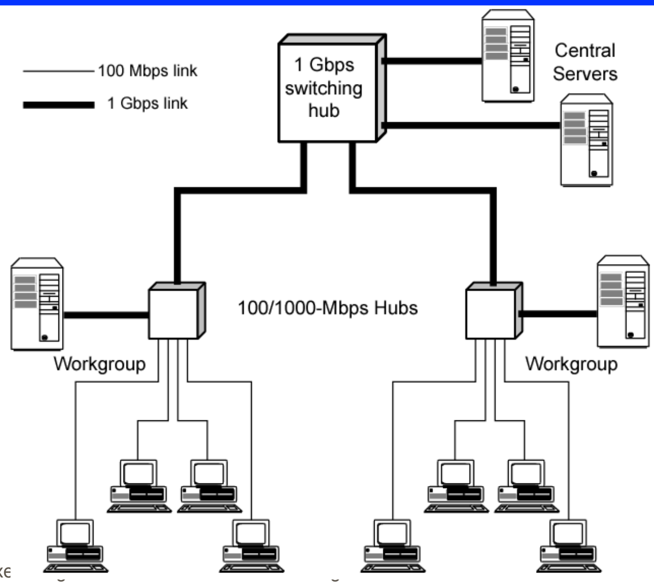

Gigabit Ethernet - Differences

- Carrier extension
- At least 4096 bit-times long
  - (for 10/100 Mbps minimum is 512)
- Frame bursting
  - allow multiple frames to be transmitted consecutively, i.e. “together”

Gigabit Ethernet – Physical

- 1000Base-SX
  - —Short wavelength, multimode fiber
- 1000Base-LX
  - —Long wavelength, Multi or single mode fiber
- 1000Base-CX
  - —Copper jumpers <25m, shielded twisted pair
- 1000Base-T
  - —4 pairs, cat 5 UTP

Gbit Ethernet Medium Options (log scale)

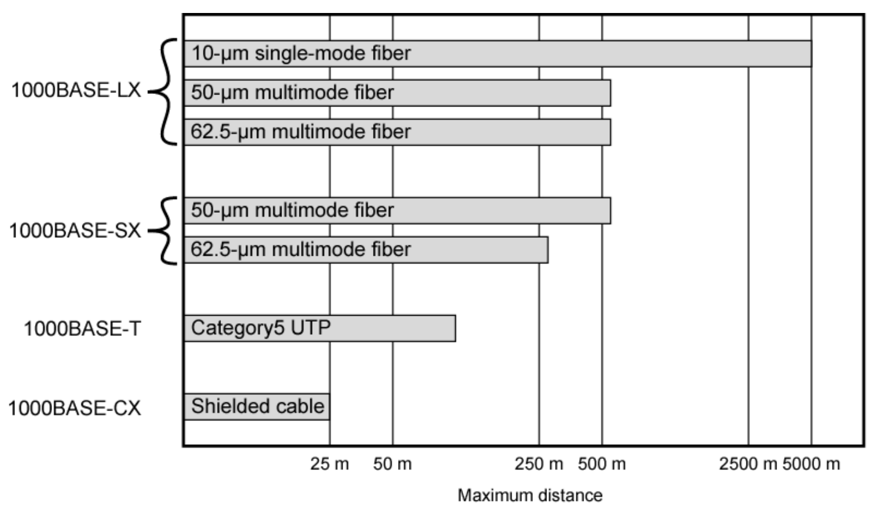

10Gbps Ethernet - Uses

- High-speed, local backbone interconnection between large-capacity switches
- Server farm, Campus wide connectivity
- Enables Internet service providers (ISPs) and network service providers (NSPs) to create very high-speed links at very low cost
- Allows construction of (MANs) and WANs
  - Connect geographically dispersed LANs between campuses or points of presence (PoPs)
- Ethernet competes with ATM and other WAN technologies
- 10-Gbps Ethernet provides substantial value over ATM

10Gbps Ethernet - Advantages

- No expensive, bandwidth-consuming conversion between Ethernet packets and ATM cells
  - Network is Ethernet, end to end
- IP and Ethernet together offer QoS and traffic policing capabilities approach those provided by ATM
- Variety of standard optical interfaces (wavelengths and link distances) specified for 10 Gb Ethernet
  - Optimizing operation and cost for LAN, MAN, or WAN

10Gbps Ethernet - Advantages

- Maximum link distances cover 300 m to 40 km
- Full-duplex mode only
- 10GBASE-S (short):
  - 850 nm on multimode fiber
  - Up to 300 m
- 10GBASE-L (long)
  - 1310 nm on single-mode fiber
  - Up to 10 km
- 10GBASE-E (extended)
  - 1550 nm on single-mode fiber
  - Up to 40 km
- 10GBASE-LX4:
  - 1310 nm on single-mode or multimode fiber
  - Up to 10 km
  - Wavelength-division multiplexing (WDM) bit stream across four light waves

10Gbps Ethernet Distance Options (log scale)

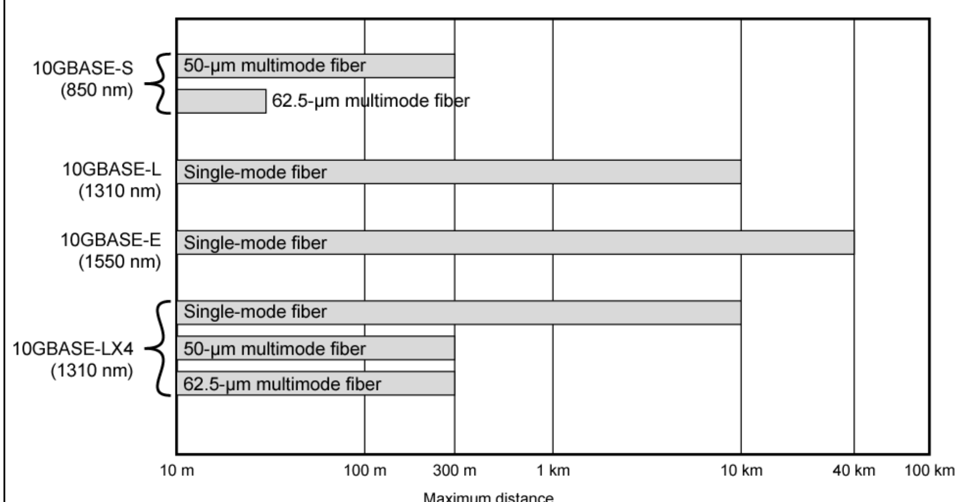

Token Ring (802.5)

- One should at least know the basic issues

Token Ring (802.5)

- Developed from IBM's commercial token ring
- Because of IBM's presence, token ring has gained broad acceptance
- Never achieved popularity of Ethernet
- Currently, there is still large installed base of token ring products
  - —Market share likely to decline

Ring Operation

- Each repeater connects to two others via unidirectional transmission links
- Single closed path
- Data transferred bit by bit from one repeater to the next
- Repeater regenerates and retransmits each bit
- Repeater performs data insertion, data reception, data removal
- Repeater acts as attachment point
- Packet removed by transmitter after one trip round ring

Listen State Functions

- Scan passing bit stream for patterns
  - —Address of attached station
  - —Token permission to transmit
- Copy incoming bit and send to attached station
  - —Whilst forwarding each bit
- Modify bit as it passes
  - —e.g. to indicate a packet has been copied (ACK)

Transmit State Functions

- Station has data
- Repeater has permission
- May receive incoming bits
  - —If ring bit length shorter than packet
    - Pass back to station for checking (ACK)
  - —May be more than one packet on ring
    - Buffer for retransmission later

Bypass State

- Signals propagate past repeater with no delay (other than propagation delay)
- Partial solution to reliability problem (see later)
- Improved performance

Ring Repeater States

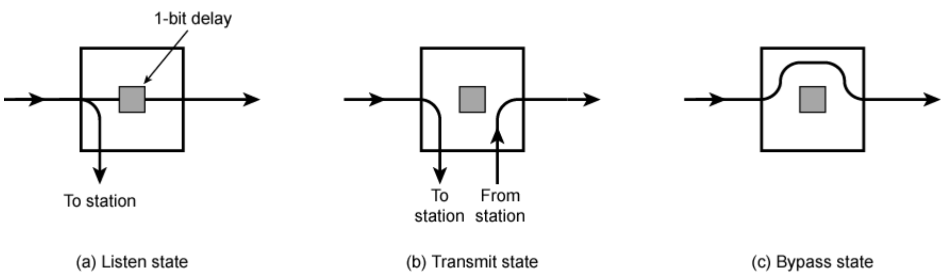

802.5 MAC Protocol

- Small frame (token) circulates when idle
- Station waits for token
- Changes one bit in token to make it SOF for data frame
- Append rest of data frame
- Frame makes round trip and is absorbed by transmitting station
- Station then inserts new token when transmission has finished and leading edge of returning frame arrives
- Under light loads, some inefficiency
- Under heavy loads, round robin

Token Ring Operation

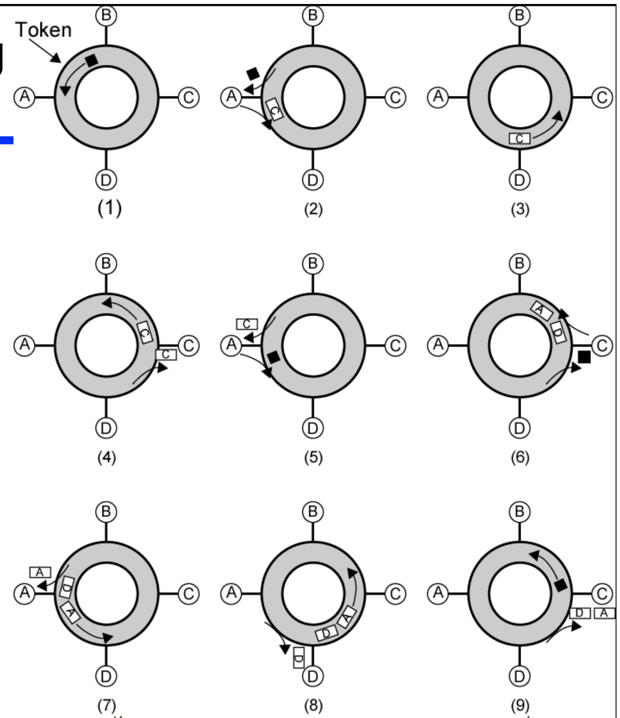

Dedicated Token Ring

- Central hub
- Acts as switch
- Full duplex point to point link
- Concentrator acts as frame level repeater
- No token passing

802.5 Physical Layer

Date rate	4	16	100	100	1000
transmission medium	UTP or STP or fiber	UTP or STP or fiber	UTP or STP	Fiber	Fiber
Signaling	Diff Manchester	Diff. Manchester	MLT-3	4B5B, NRZI	8B/10B
Max frame size	4550	18,200	18,200	18,200	18,200
Access control	TP or DTR	TP or DTR	DTR	DTR	DTR

TP = token passing access control
DTR = dedicated token ring

Fibre Channel - Background

- I/O channel
  - Direct point to point or multipoint communication link
  - Hardware based
  - High Speed
  - Very short distance
  - User data moved from source buffer to destination buffer
- Network connection
  - Interconnected access points
  - Software based protocol
  - Flow control, error detection &recovery
  - End systems connections

Fibre Channel

- Best of both technologies
- Channel oriented
  - —Data type qualifiers for routing frame payload
  - —Link level constructs associated with I/O ops
  - —Protocol interface specifications to support existing I/O architectures
    - e.g. SCSI
- Network oriented
  - —Full multiplexing between multiple destinations
  - —Peer to peer connectivity
  - —Internetworking to other connection technologies

Fibre Channel Requirements

- Full duplex links with two fibers per link
- 100 Mbps to 800 Mbps on single line
  - — Full duplex 200 Mbps to 1600 Mbps per link
- Up to 10 km
- Small connectors
- High-capacity utilization, distance insensitivity
- Greater connectivity than existing multidrop channels
- Broad availability
  - — i.e. standard components
- Multiple cost/performance levels
  - — Small systems to supercomputers
- Carry multiple existing interface command sets for existing channel and network protocols
- Uses generic transport mechanism based on point-to-point links and a switching network
- Supports simple encoding and framing scheme
- In turn supports a variety of channel and network protocols

Fibre Channel Elements

- End systems - Nodes
- Switched elements - the network or fabric
- Communication across point to point links

Fibre Channel Network

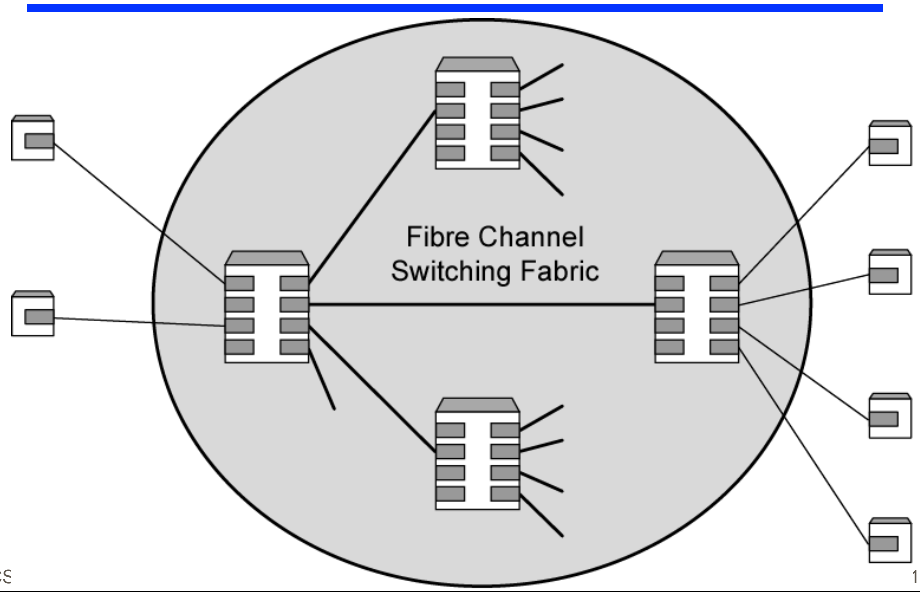

Fibre Channel Protocol Architecture (1)

- FC-0 Physical Media
  - Optical fiber for long distance
  - coaxial cable for high speed short distance
  - STP for lower speed short distance
- FC-1 Transmission Protocol
  - 8B/10B signal encoding
- FC-2 Framing Protocol
  - Topologies
  - Framing formats
  - Flow and error control
  - Sequences and exchanges (logical grouping of frames)

Fibre Channel Protocol Architecture (2)

- FC-3 Common Services
  - Including multicasting
- FC-4 Mapping
  - Mapping of channel and network services onto fibre channel
    - e.g. IEEE 802, ATM, IP, SCSI

Fibre Channel Physical Media

	800 Mbps	400 Mbps	200 Mbps	100 Mbps
Single mode fiber	10 km	10 km	10 km	—
50-μm multimode fiber	0.5 km	1 km	2 km	—
62.5-μm multimode fiber	175 m	1 km	1 km	—
Video coaxial cable	50 m	71 m	100 m	100 m
Miniature coaxial cable	14 m	19 m	28 m	42 m
Shielded twisted pair	28 m	46 m	57 m	80 m

Fibre Channel Fabric

- General topology called fabric or switched topology
- Arbitrary topology includes at least one switch to interconnect number of end systems
- May also consist of switched network
  - Some of these switches supporting end nodes
- Routing transparent to nodes
  - Each port has unique address
  - When data transmitted into fabric, edge switch to which node attached uses destination port address to determine location
  - Either deliver frame to node attached to same switch or transfers frame to adjacent switch to begin routing to remote destination

Fabric Advantages

- Scalability of capacity
  - As additional ports added, aggregate capacity of network increases
  - Minimizes congestion and contention
  - Increases throughput
- Protocol independent
- Distance insensitive
- Switch and transmission link technologies may change without affecting overall configuration
- Burden on nodes minimized
  - Fibre Channel node responsible for managing point-to-point connection between itself and fabric
  - Fabric responsible for routing and error detection

Alternative Topologies

- Point-to-point topology
  - Only two ports
  - Directly connected, with no intervening switches
  - No routing
- Arbitrated loop topology
  - Simple, low-cost topology
  - Up to 126 nodes in loop
  - Operates roughly equivalent to token ring
- Topologies, transmission media, and data rates may be combined

Five Applications of Fibre Channel

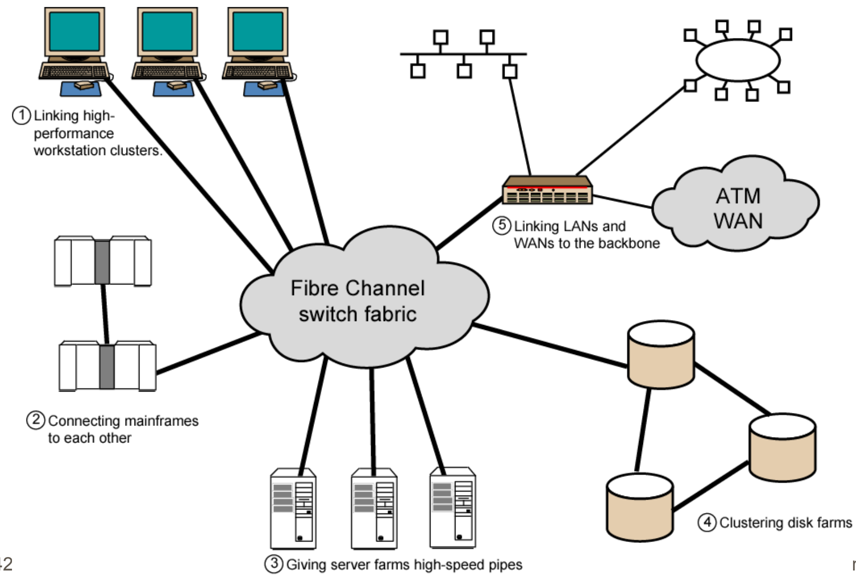

Fibre Channel Prospects

- Backed by Fibre Channel Association
- Interface cards for different applications available
- Most widely accepted as peripheral device interconnect
  - To replace such schemes as SCSI
- Technically attractive to general high-speed LAN requirements
- Must compete with Ethernet and ATM LANs
- Cost and performance issues should dominate the consideration of these competing technologies

Summary

- High speed LANs emergence
- Ethernet technologies
  - —CSMA & CSMA/CD media access
  - —10Mbps ethernet
  - —100Mbps ethernet
  - —1Gbps ethernet
  - —10Gbps ethernet
- Fibre Channel