Fast Ethernet Training

Hadriel Kaplan
InterOperability Lab
UNH

Topics

- A brief history of Ethernet and the IEEE
- What is Fast Ethernet?
- Fast Ethernet Internals
- Fast Ethernet network design constraints
- Why use it, and when?
- How to buy Fast Ethernet equipment (what to look for)

Topics

- **A (not-so) brief history of Ethernet and the IEEE**
  - where it fits in the OSI stack
  - the IEEE organization
  - Aloha and slotted Aloha protocols
  - Robert Metcalf et al
  - Ye Olde Ethernet

The ISO OSI Stack Model

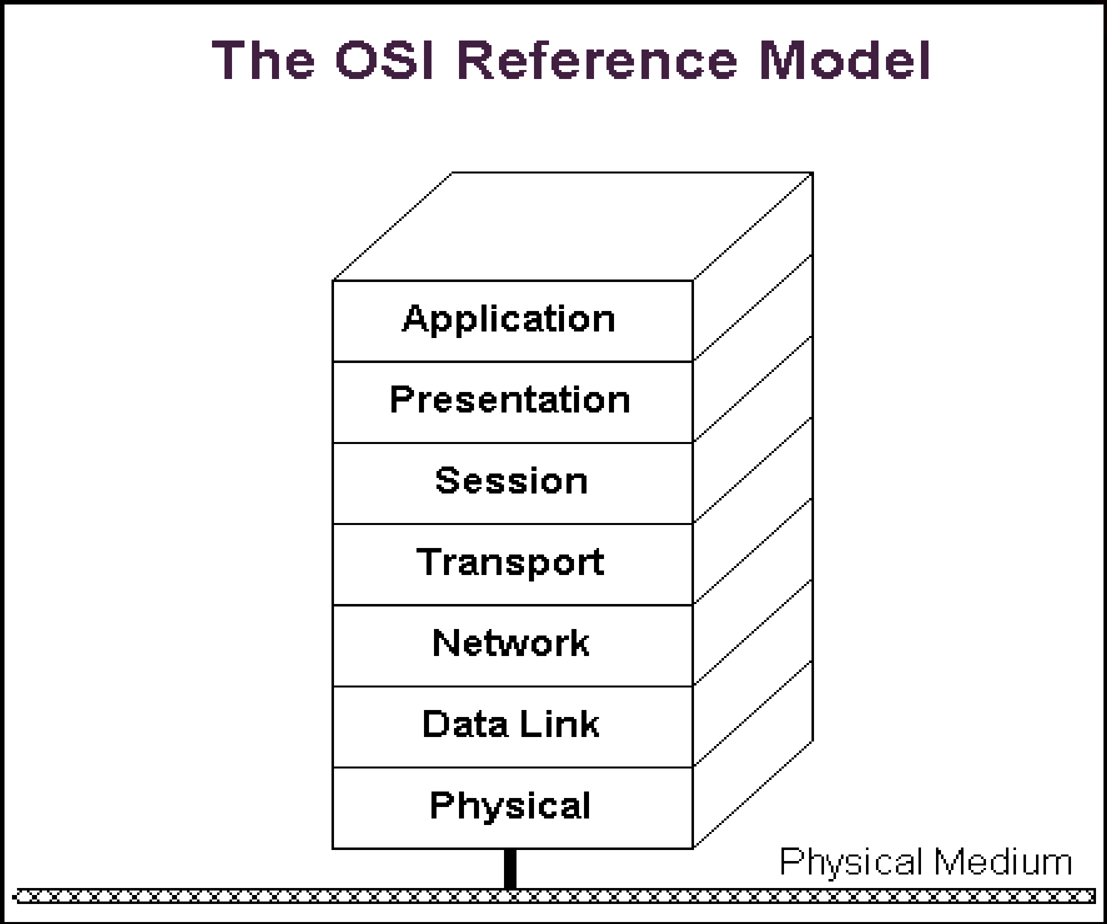

- IEEE specifies LAN/MAN physical and data link layers
- IETF specifies upper layers
- ANSI specs physical mediums

IEEE Relationship to OSI

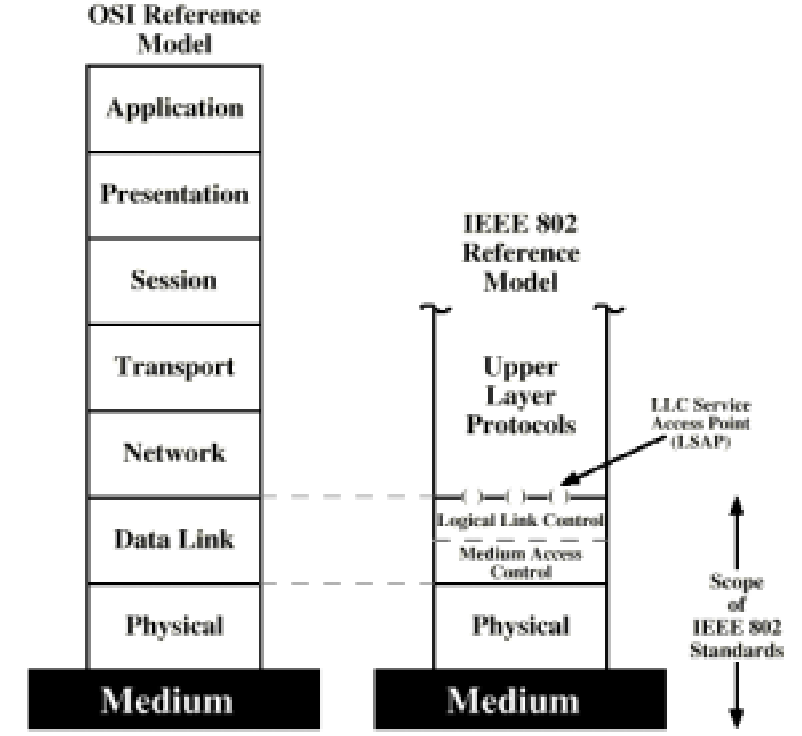

Figure 12.1 IEEE 802 Protocol Layers Compared to OSI Model

IEEE Working Groups

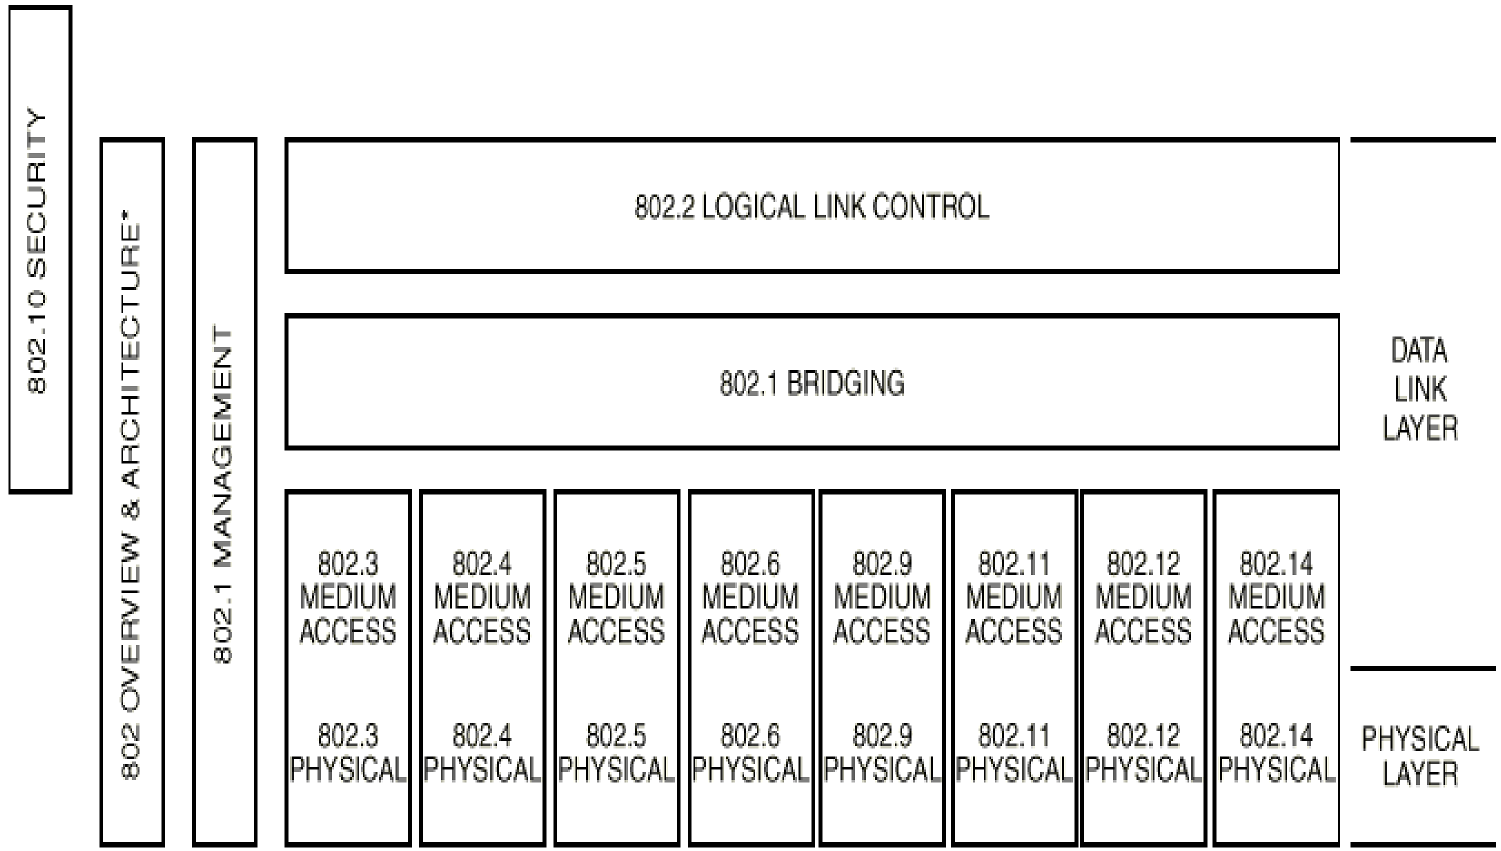

* Formerly IEEE Std 802.1A.

The IEEE 802.3 Committee

IEEE Ethernet Standards/Working Groups

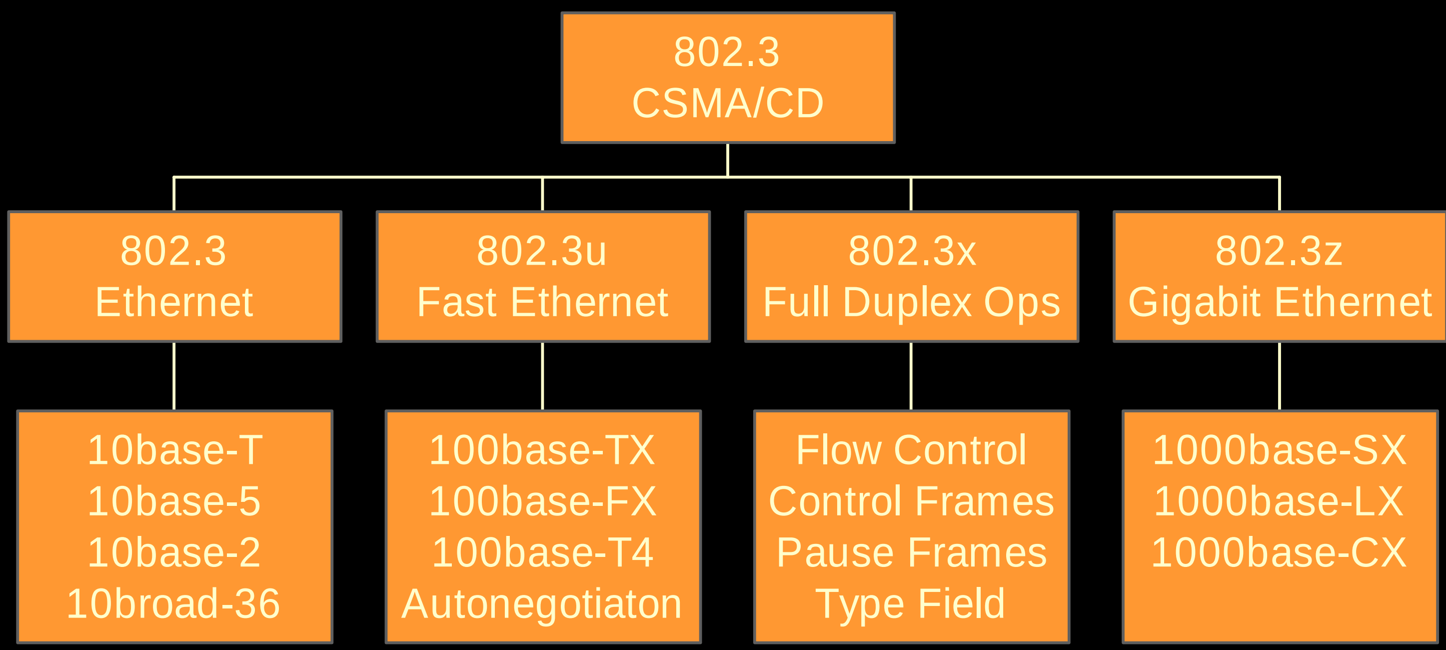

Pure Aloha

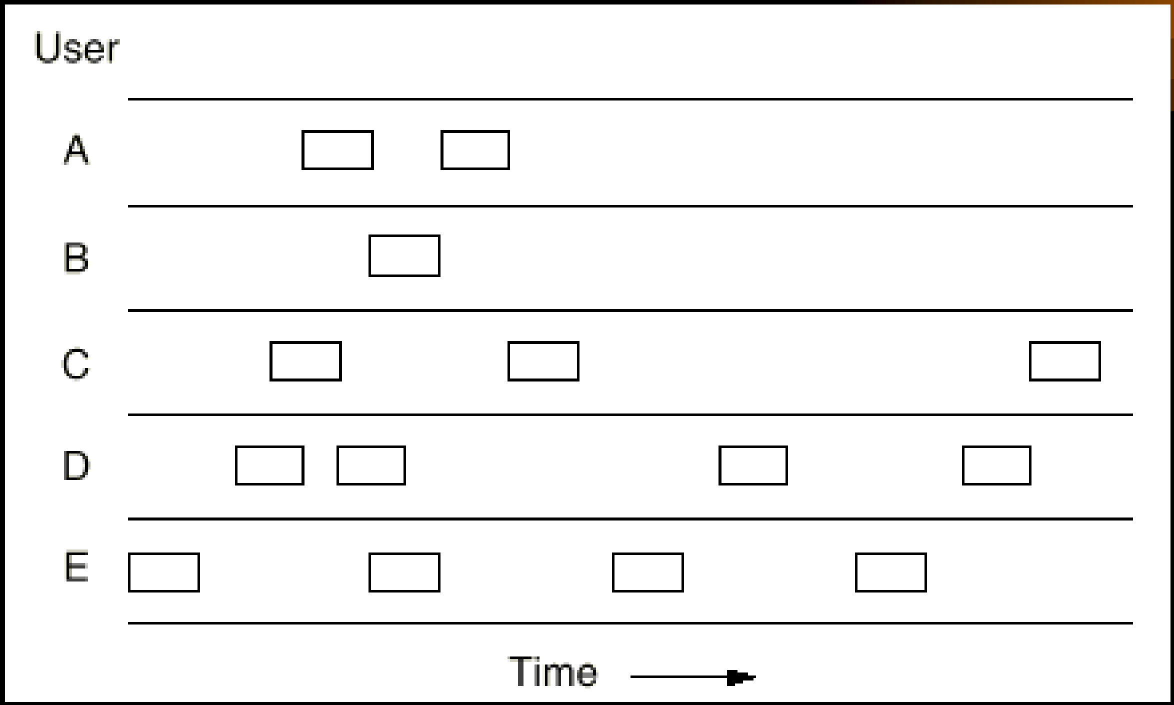

- Transmit when you want to, regardless of others.

Pure Aloha Collisions

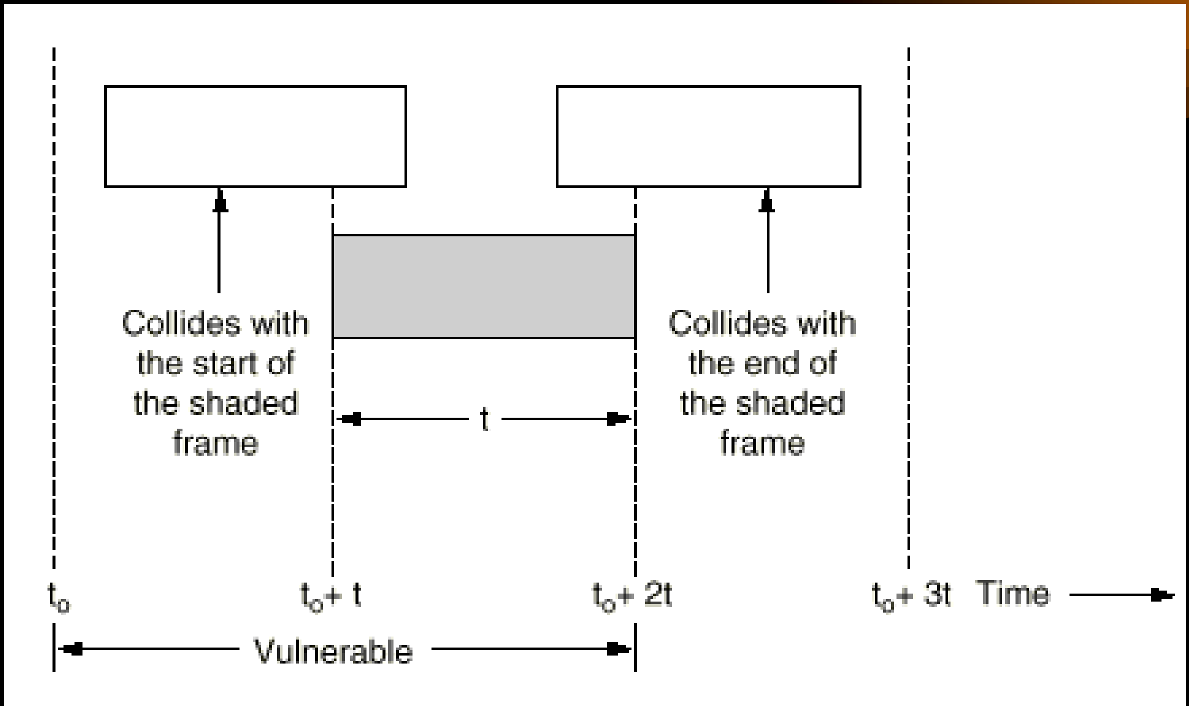

- Extremely inefficient, since the worst-case period of vulnerability is the time to transmit two frames.

Slotted Aloha

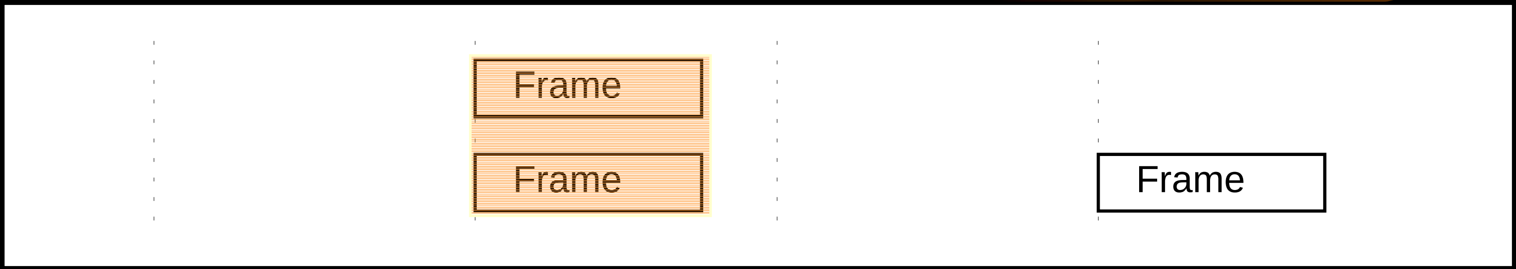

- Transmit only at the beginning of synchronized "slot times"
- Collision inefficiency limited to one frame transmission time

Aloha vs. Slotted Aloha

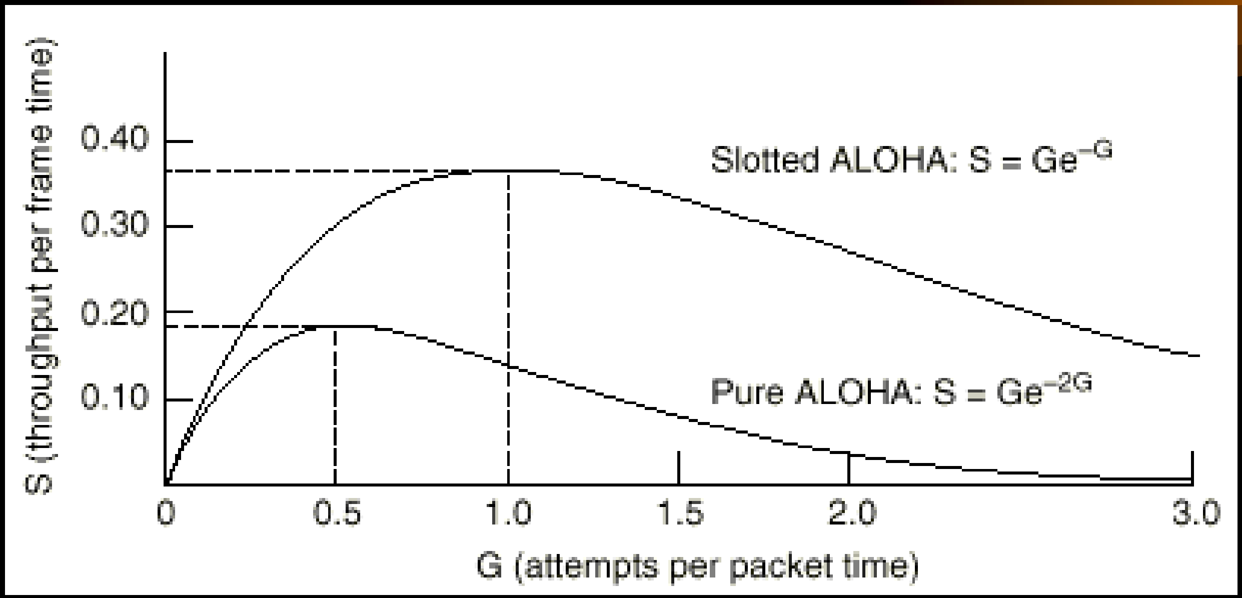

- Throughput efficiency increases dramatically for Slotted Aloha.

Robert Metcalf's Idea

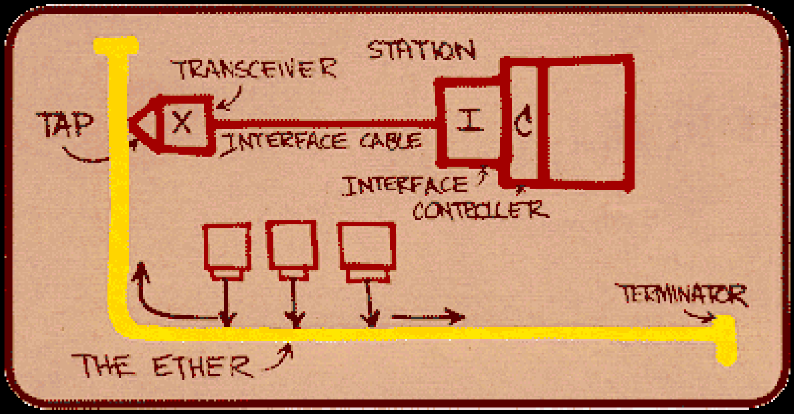

- Invented by Metcalf at Xerox in 1973 and patented in 1976
- Xerox convinced Digital and Intel to join in making products (hence the group called DIX)
- IEEE standard in 1989

CSMA/CD Defined

CS - Carrier Sense (Is someone already talking?)

MA - Multiple Access (I hear what you hear!)

CD - Collision Detection (Hey, we're both talking!)

1. If the medium is idle, transmit anytime.
2. If the medium is busy, wait and transmit right after.
3. If a collision occurs, backoff for a random period, then go back to 1.

We use CSMA/CD in normal group conversation.

CSMA/CD

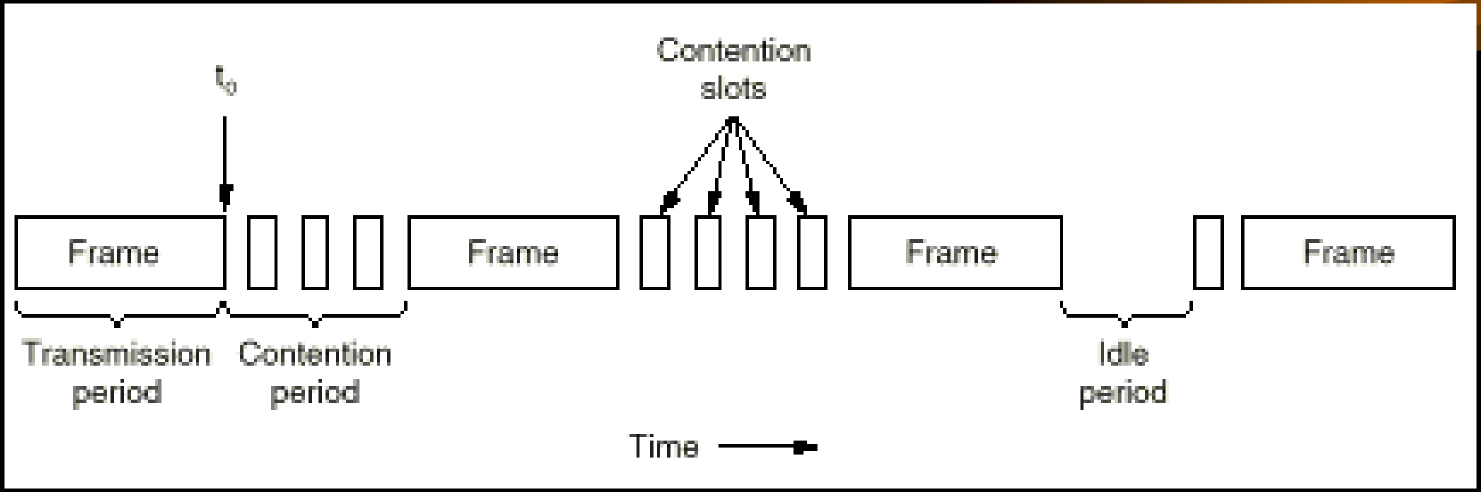

- CSMA/CD can be in one of three states: contention, transmission, or idle.

Plain Vanilla 802.3

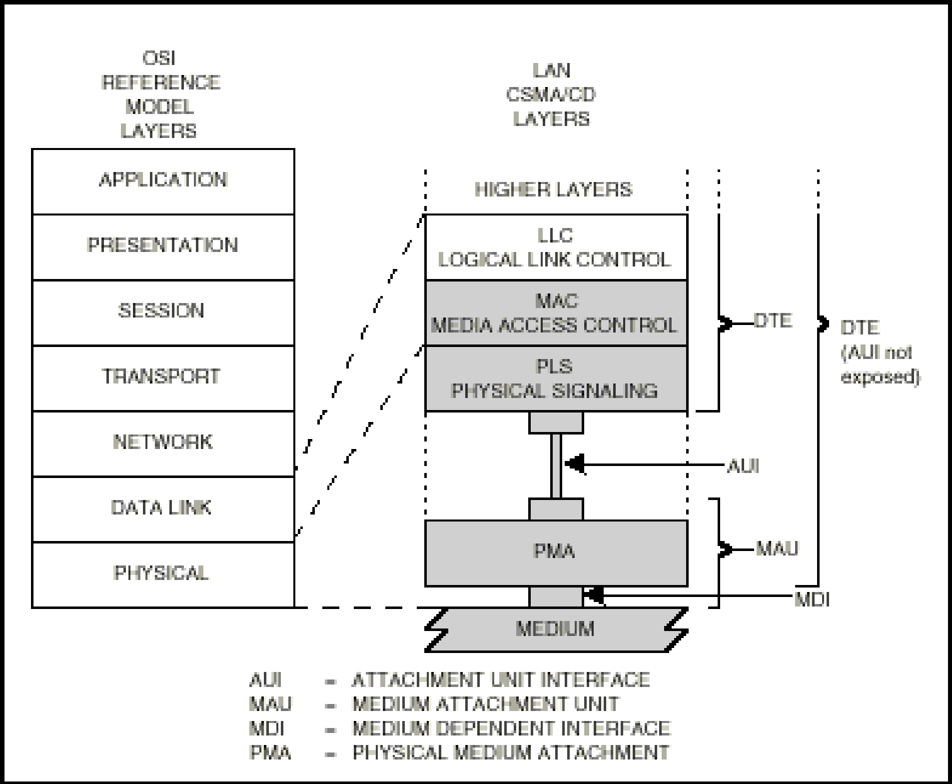

- Defined MAC
- Defined PHY
  - 10base-5
  - 10base-2
  - 10base-T
  - 10base-F
- Defined Repeater

Ethernet Frame Format

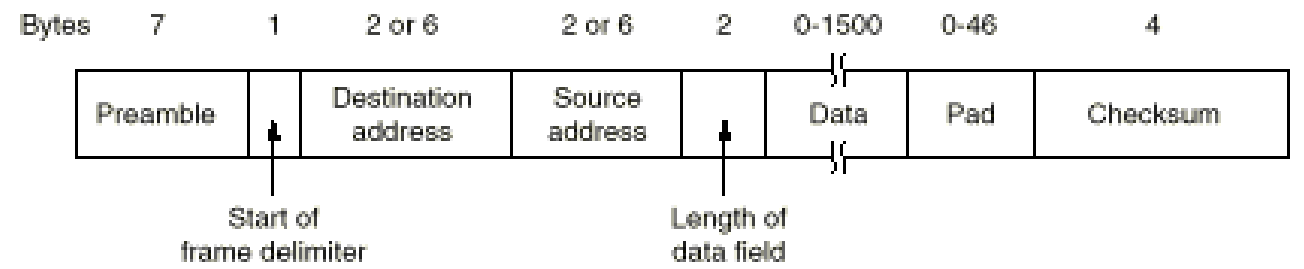

Fig. 4-21. The 802.3 frame format.

- 802.3 Ethernet Frame Format:
  - Length field for length of data (not pad)
  - 7 bytes preamble, 1 byte SFD
- Ethernet type II Frame Format (DIX):
  - Type field for what layer data belongs to (including pad)
  - 8 bytes Preamble, but it looks the same as 802.3 format

Ethernet Flavors

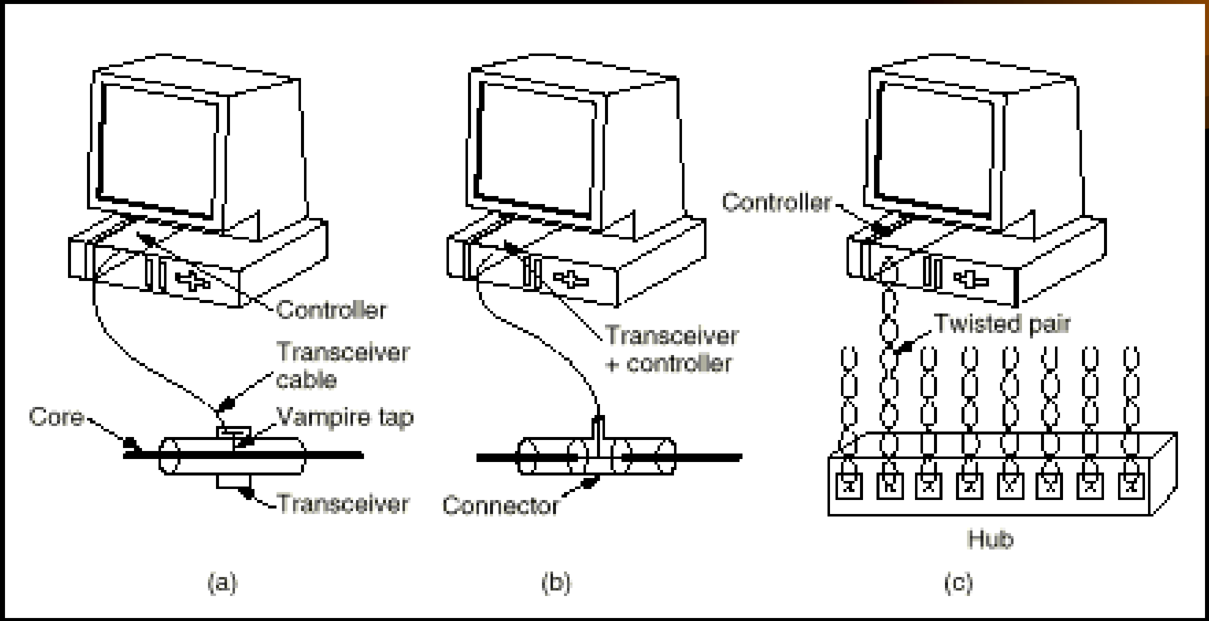

- Because 10base-T devices always linked to the repeater, it must make the network appear to be shared, as if it were 10base-5 or 10base-2

Repeaters

- Works at layer 1 (PHY layer) ONLY
  - thus it doesn't understand frame formats
- Repeat incoming signal from a port to all other ports with:
  - restored timing
  - restored waveform shape
  - very little delay
- If 2 or more simultaneous receptions, transmit jam
- Can connect dissimilar media/PHY types (e.g., 10base-T and 10base-2)

Topics

- A brief history of Ethernet and the IEEE
- **What is Fast Ethernet?**
  - Same Ethernet Frame Format
  - What makes it *Fast* Ethernet?

Ethernet Frame Format

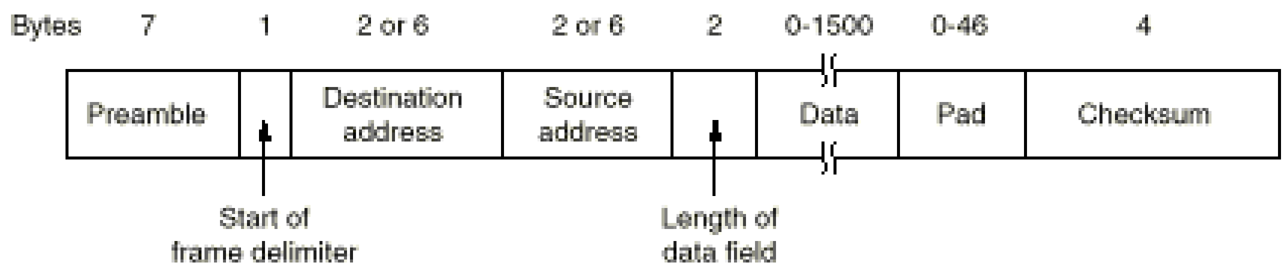

Fig. 4-21. The 802.3 frame format.

- Same old frame format
- In fact, same MAC layer (except runs faster)
- Of course, the same frame now takes 1/10 the time to send!

What makes it Fast *Ethernet?*

- It runs at 100mbps (data rate)
- Still uses 2-pairs of UTP
- Uses same frame types, lengths, formats
- Still on a shared network (half duplex)
- Still CSMA/CD
- Still dirt cheap and trivial to install

Topics

- A brief history of Ethernet and the IEEE
- What is Fast Ethernet?
- **Fast Ethernet Internals**
  - How it compares to 10base-T
  - Encoding schemes
  - Line signaling
  - The many flavors of Fast Ethernet

How it stacks up with 10base

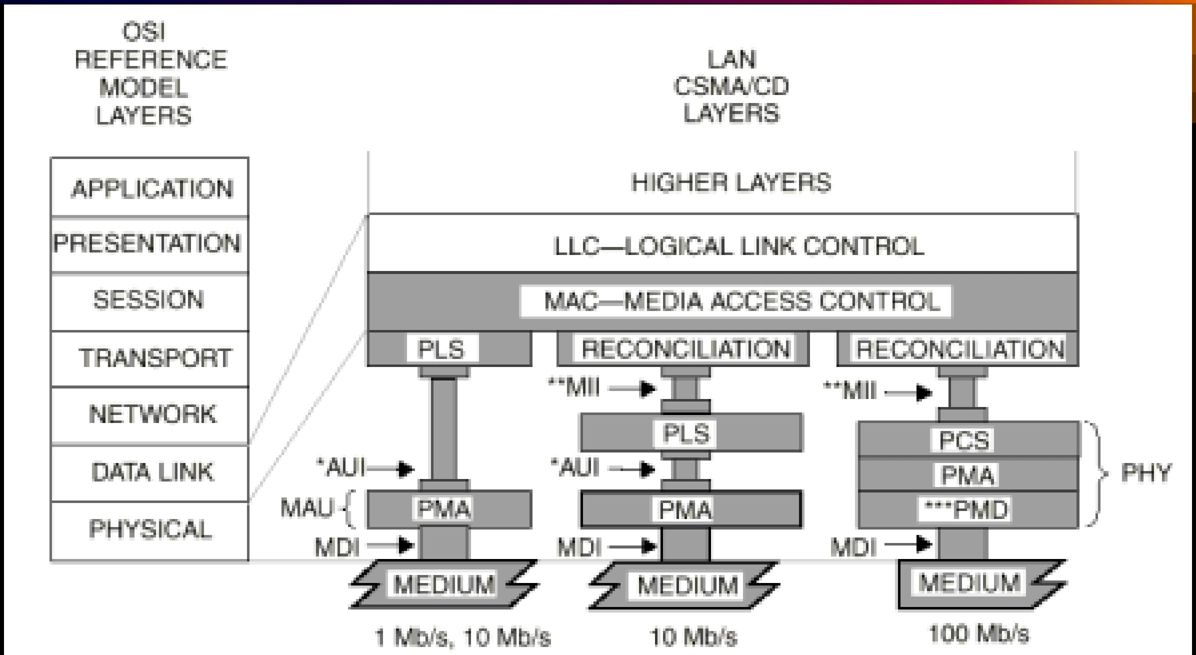

- Same MAC, whole new PHY (courtesy of FDDI).
- More layers defined to allow for easier swapping/interfacing of components

Cable Types

	Attenuation (dB per 100 m)			Near-end Crosstalk (dB)		
Frequency (MHz)	Category 3 UTP	Category 5 UTP	150 Ω STP	Category 3 UTP	Category 5 UTP	150 Ω STP
1	2.6	2.0	1.1	41	62	58
4	5.6	4.1	2.2	32	53	58
16	13.1	8.2	4.4	23	44	50.4
25	—	10.4	6.2	—	32	47.5
100	—	22.0	12.3	—	—	38.5
300	—	—	21.4	—	—	31.3

- 10base-T runs on Category 3 UTP or higher
- 100base-TX runs only on Category 5 UTP cable
- 100base-T4 and T2 run on Cat 3 UTP as well
- Still uses pairs 1/2, 3/6

4B/5B Encoding

- Allows for control codes
- Adds error detection

Table 13.8 4B/5B Code Groups (page 1 of 2)

Data Input (4 bits)	Code Group (5 bits)	NRZI pattern	Interpretation
0000	11110		Data 0
0001	01001		Data 1
0010	10100		Data 2
0011	10101		Data 3
0100	01010		Data 4
0101	01011		Data 5
0110	01110		Data 6
0111	01111		Data 7
1000	10010		Data 8
1001	10011		Data 9
1010	10110		Data A
1011	10111		Data B
1100	11010		Data C

Table 13.8 4B/5B Code Groups (page 2 of 2)

1101	11011		Data D
1110	11100		Data E
1111	11101		Data F
	11111		Idle
	11000		Start of stream delimiter, part 1
	10001		Start of stream delimiter, part 2
	01101		End of stream delimiter, part 1
	00111		End of stream delimiter, part 2
	00100		Transmit error
	other		invalid codes

Speed problems

- 4B/5B encoding “wastes” data bandwidth, so we increase line rate from 100mbps to 125mbps to keep “real data” rate 100mbps
- On fiber that’s fine, but you can’t transmit 125MHz signals easily (over 22dB atten.!) or legally on UTP (that darn FCC!)
- Solution: use a forgiving line encoding scheme, and scramble the data

Manchester Encoding

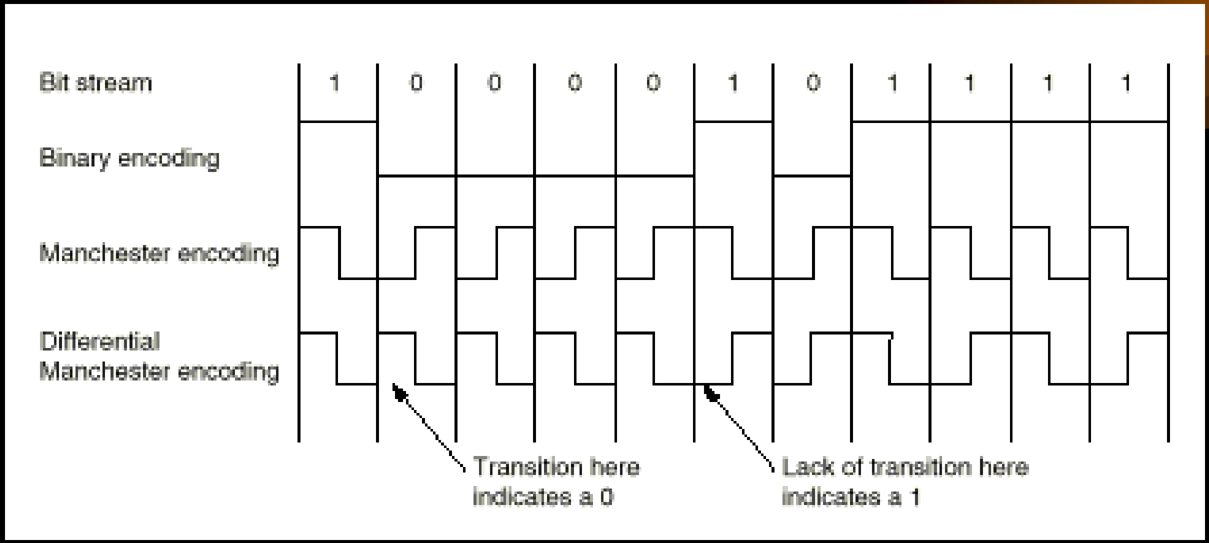

- 10base-T uses Manchester Encoding
- + to - transition = “1” - to + = “0”
- Always DC balanced - Always has a transition each bit-time for clock recovery.

MLT-3 Transmission Encoding

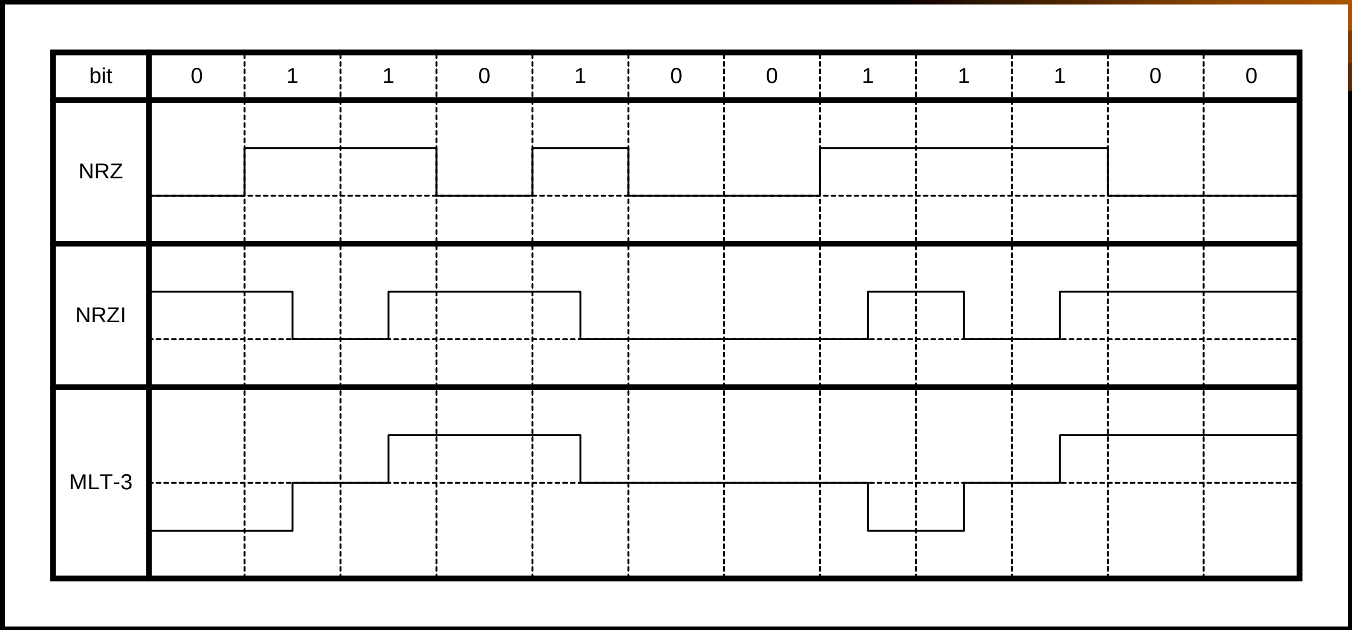

- If data is "1", then transition from current level to next level. ie 1111=+1,0,-1,0 reducing freq by 1/4
- If bit data is "0", then don't transition
- 100base-FX uses NRZ-I

Scrambled Eggs

- Even MLT-3 wasn’t quiet enough, so they scramble the bits before they MLT-3 encode them on the wire
- This means the receiving device needs to descramble the scrambled data
- That means the receiving device’s descrambler needs to be continuously synchronized
- So we send a well known control code (called IDLE) whenever we aren’t sending data (it also keeps the clock synched, so preamble is superfluous)

More Stacks!

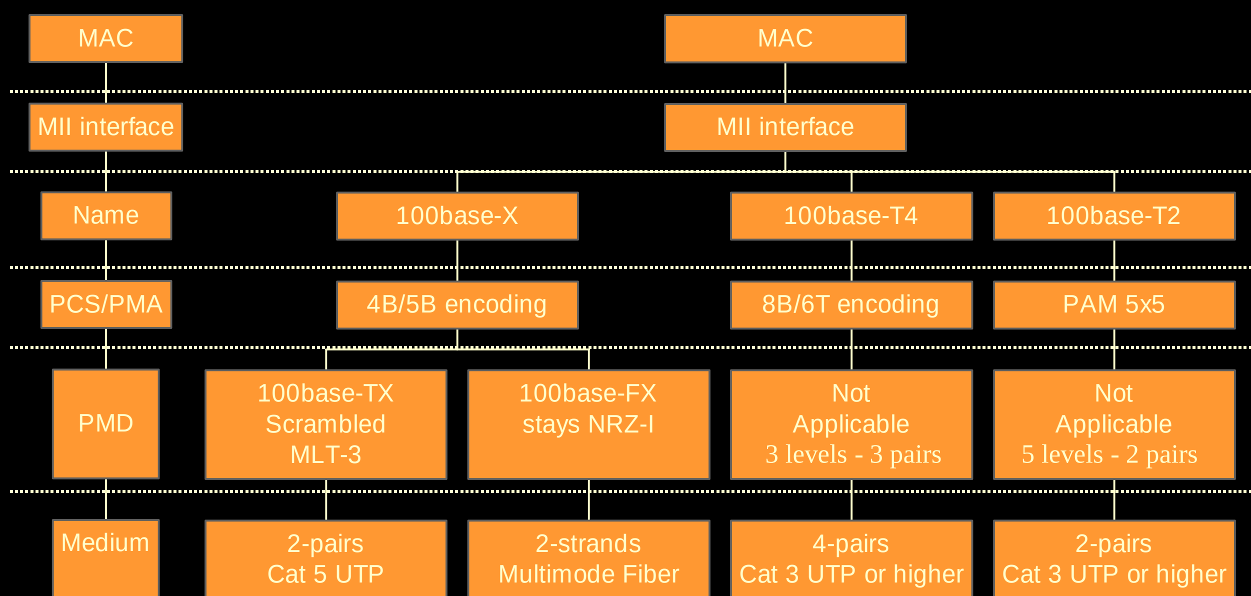
MAC
MII interface
Name
PCS/PMA
PMD
Medium
MAC
MII interface
100base-X
100base-T4
100base-T2
4B/5B encoding
8B/6T encoding
PAM 5x5
100base-TX Scrambled MLT-3
100base-FX stays NRZ-I
Not Applicable 3 levels - 3 pairs
Not Applicable 5 levels - 2 pairs
2-pairs Cat 5 UTP
2-strands Multimode Fiber
4-pairs Cat 3 UTP or higher
2-pairs Cat 3 UTP or higher

Intermission

You can wake up now...

Full Duplex - The ATM Killer

- Remember CSMA/CD? Now forget it.
- Ever since 10base-T, collisions were purely logical - there weren't really collisions on the wire, since you have separate channels for transmit/receive
- So as long as we don't need to share a network (like with a repeater), why bother "colliding"?
- New MAC: transmitting while receiving is OK.
- Still maintain IPG, frame sizes, and physical layer

Full Duplex Forever!

- Pros:
  - aggregate throughput = 200mbps
  - no collision efficiency penalty
  - <u>no collision domain</u>
- Cons:
  - must be a point-to-point link (i.e., no repeaters)
  - so darn fast it might overload the receiving side
  - no backpressure ability (except for Pause frames)

Autonegotiation vs. Autosensing

- Autonegotiation
  - standardized speed handshake
  - auto-configures to best possible link (e.g., 100 full duplex)
  - still links with older or non-autoneg devices
  - sometimes causes autosensing (NOT autonegotiating) devices to link at 10 and not 100
- Autosensing/Speed Detection
  - several different proprietary methods
  - only auto-configures to 10 or 100, not duplex settings
  - creates many interoperability headaches

Autonegotiation - How?

- Constantly sends out 10base-T Link Test Pulses before linking
- The pulses are grouped together in defined “words” that convey meanings, such as “I can do 100 half and full duplex”
- Older 10base-T devices just think they’re LTPs
- Autoneg devices understand them as words and exchange handshake info to link at best possible link
- If the autoneg device sees regular LTPs coming in (not autoneg words), it just links at 10 half duplex
- If the autoneg device sees Fast Ethernet IDLE stream coming in, it just links at 100 half duplex.
- Unfortunately, this only works on copper

Topics

- A brief history of Ethernet and the IEEE
- What is Fast Ethernet?
- Fast Ethernet Internals
- **Fast Ethernet network design constraints**
  - physical link limit
  - collision domain
  - full-duplex
  - users and losers

Physical Link Limits

- Attenuation reduces signal amplitude, so 100base-TX can only run ~100 meters before the signal must be repeated.
- 100base-FX can go a *long* way (2-5 km)

The Collision Domain

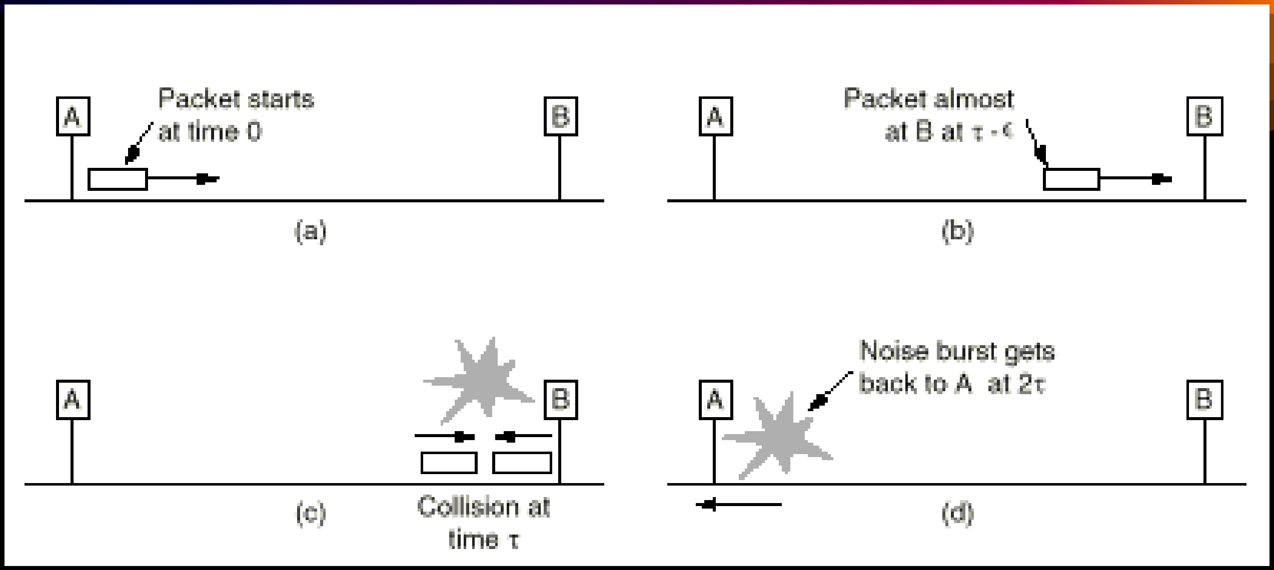

- Collision detection can take as long as 2t, worst case.
- This “round-trip” delay defines the max Ethernet network diameter, or collision domain.
- Round-trip delay = 512 bit times for all ethernets.

Collision Domain for 10base

- If you had a perfect cable, you could run longer than 4km and still be within the collision domain.
- Cable is not perfect. Copper’s slower than fiber, and repeaters are *really* slow.
- The 10base-2 and T rule of thumb for farthest end to farthest end of a domain is 5-4-3:
  - 5 total segments
  - 4 repeaters
  - 3 populated (if 10base-2)

Serious Limit

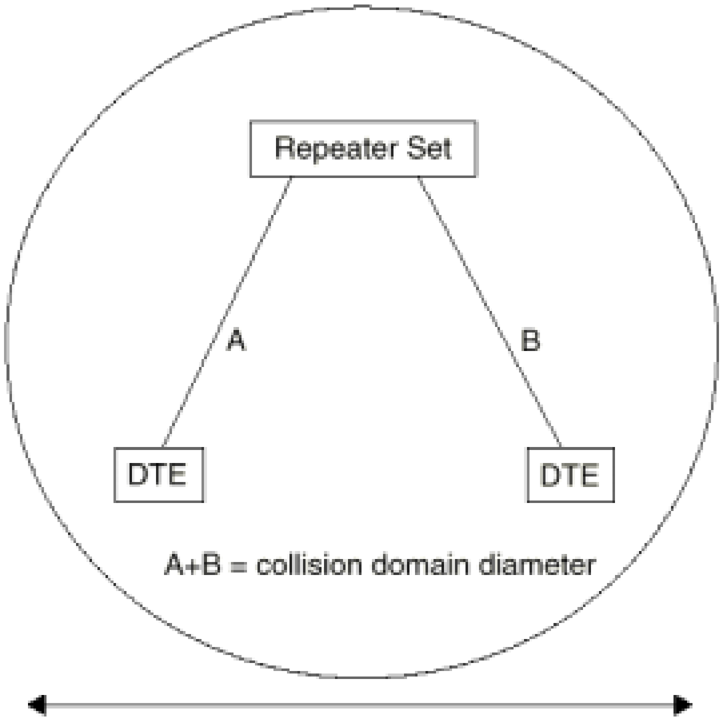

See table 29-2 for maximum collision domain diameter.

Figure 29-4—Model 1: Single repeater

Table 29-2—Maximum Model 1 collision domain diameter[a]

Model	Balanced cable (copper)	Fiber	Balanced cable & fiber (T4 and FX)	Balanced cable & fiber (TX and FX)
DTE-DTE (see figure 29-3)	100	412	na	na
One Class I repeater (see figure 29-4)	200	272	231[b]	260.8[b]
One Class II repeater (see figure 29-4)	200	320	304[b,c]	308.8[b]
Two Class II repeaters (see figure 29-5)	205	228	236.3[d,c]	216.2[d]

[a]In meters, no margin.
[b]Assumes 100 m of balanced cable and one fiber link.
[c]This entry included for completeness. It may be impractical to construct a T4 to FX class II repeater.
[d]Assumes 105 m of balanced cable and one fiber link.

- 512 bit times isn't much for F.E., because the bit time is 1/10 what it was for 10mbps
  - Even on fiber, the max diameter is 412 meters, and that's purely because of the round-trip time.

Repeater Types

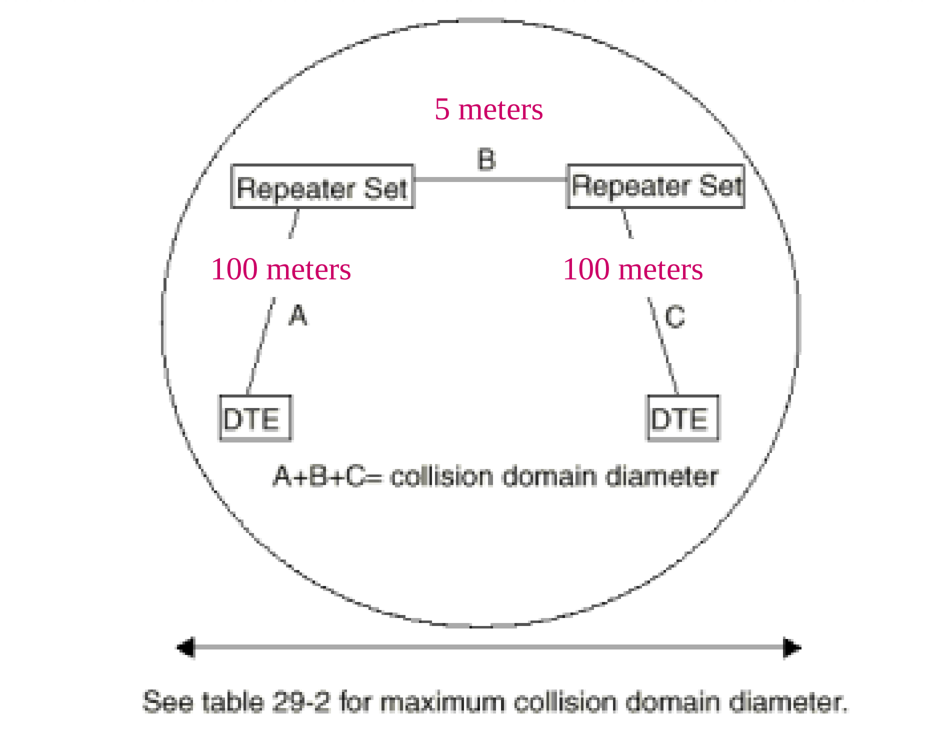

Figure 29-5—System Model 1: Two Class II repeaters

- Repeater delay is VERY significant. So much so, they defined two types or speeds of repeaters:
  - Type I are slower
  - Type II are faster
  - Even using a Type II, you can only have 2 of them in a network!

Breaking up Collision Domains

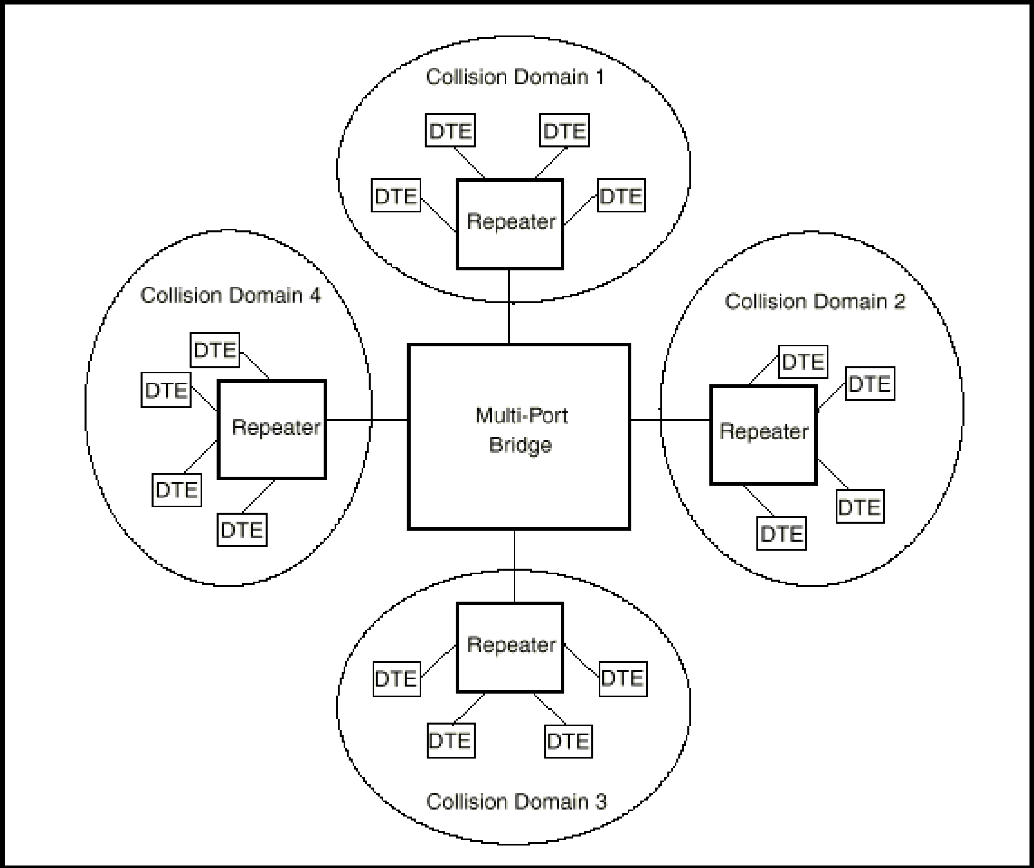

- Repeaters are inside the collision domain, since they propagate collisions
- Bridges/Switches break up the domains, since they operate at layer 2 and buffer packets before sending them

Remember Full Duplex?

- One major benefit is **no collision domain**, since there are no collisions!
- For 100base-TX, 100 meters is still a limitation (because it was a physical one)
- But 100base-FX can go over 2km when the MAC layer is set to full-duplex, because the limitation is physical (attenuation)

Saturating the Slower Segments

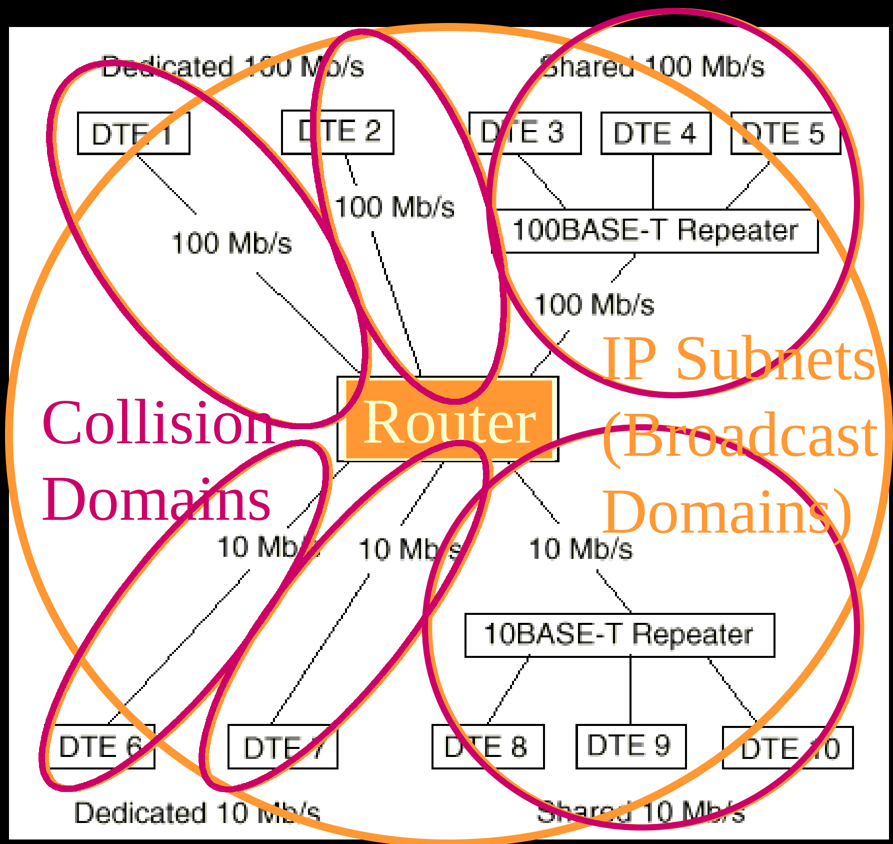

- In a bridged network, broadcast and multicast traffic is sent everywhere
- 100mbps traffic could thus congest 10mbps networks
- Using VLANs or separate IP subnets (with routers) is the only solution

Topics

- A brief history of Ethernet and the IEEE
- What is Fast Ethernet?
- Fast Ethernet Internals
- Fast Ethernet network design constraints
- **Why use it, and when?**

Why use it and when?

- Why:
  - fast (duh)
  - cheap (barely a price difference)
  - easy (same rules, mostly)
- When:
  - when physical issues are not a problem (Cat 5 UTP, 100 meters limit, collision diameter, etc.)
  - when 10mbps users aren’t going to be saturated by 100mbps traffic (broadcast and multicast)

Topics

- A brief history of Ethernet and the IEEE
- What is Fast Ethernet?
- Fast Ethernet Internals
- Fast Ethernet network design constraints
- Why use it, and when?
- **How to buy Fast Ethernet equipment (what to look for)**

Future Standards

- 802.3 CSMA/CD
  - 802.3 Ethernet
    - 10base-T
    - 10base-5
    - 10base-2
    - 10broad-36
  - 802.3u Fast Ethernet
    - 100base-TX
    - 100base-FX
    - 100base-T4
    - Autonegotiaton
  - 802.3x Full Duplex Ops
    - Flow Control
    - Control Frames
    - Pause Frames
    - Type Frames
  - 802.3z Gigabit Ethernet
    - 1000base-SX
    - 1000base-LX
    - 1000base-CX
  - 802.3ab Gig over UTP
    - 1000base-T
  - 802.3ac Frame Tagging
    - VLAN Tags
  - 802.3ad Link Aggregation