Chapter 3
Physical and Link Layer 1:
Ethernet

Ethernet

- The most popular physical layer implementation
- Developed in by 1982, DEC, Xerox and Intel
- Uses *Carrier sense, multiple access with collision detection (CSMA/CD)*
- Initially a 10Mb/s signalling rate
- Nearly the same as IEEE 802.3

Ethernet

- An Ethernet *frame*:

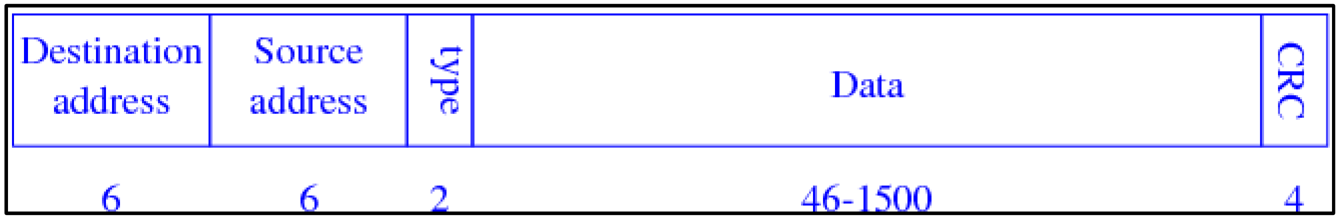

- 2 byte *type* that indicates what kind of data follows, e.g., 0800 for an IP packet
- Then the data, maximum 1500 bytes, *minimum 46 bytes*
- Data field must be padded with extra bytes if fewer than 46 bytes are supplied

Ethernet

- An Ethernet *frame*:

Destination address	Source address	type	Data	CRC
6	6	2	46-1500	4

- 4 byte *checksum,* also called *cyclic redundancy check (CRC)*
- Used to check for errors in the frame

Ethernet

- Ethernet is a *shared medium*

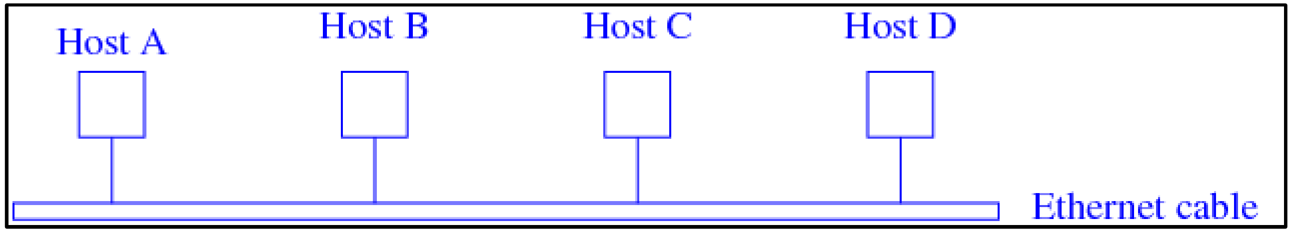

- A wishes to send to B
- If C is already sending to D, then A must wait

Ethernet

- Ethernet is a *shared medium*

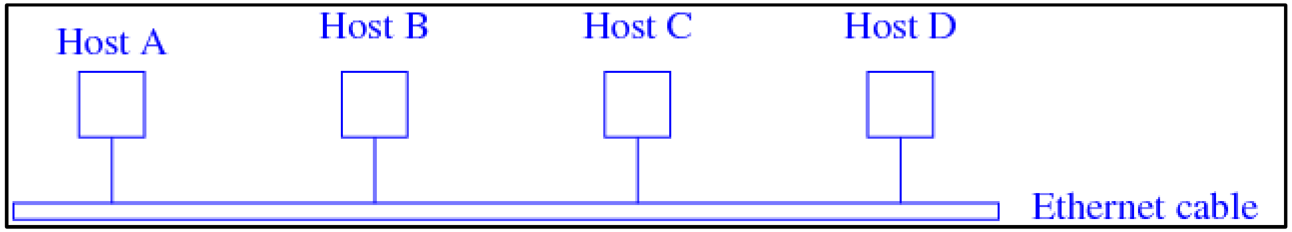

- A sends packet, but keeps listening for a clash (*collision detection*)

Ethernet

- Ethernet is a *shared medium*

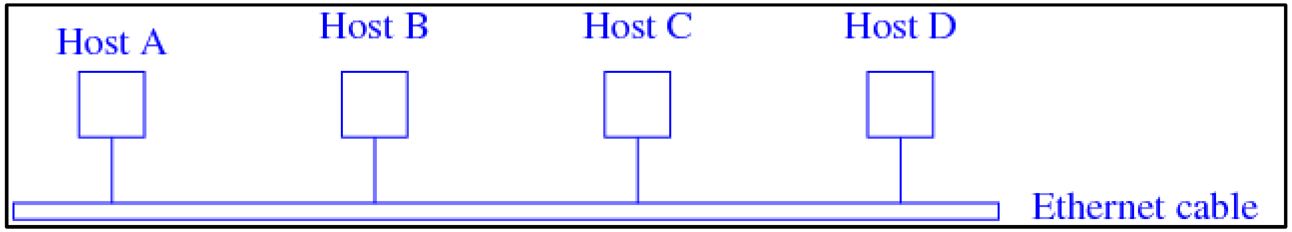

- Random wait means one of A or C gets in next, the other sees this through its carrier sense

Ethernet Hardware

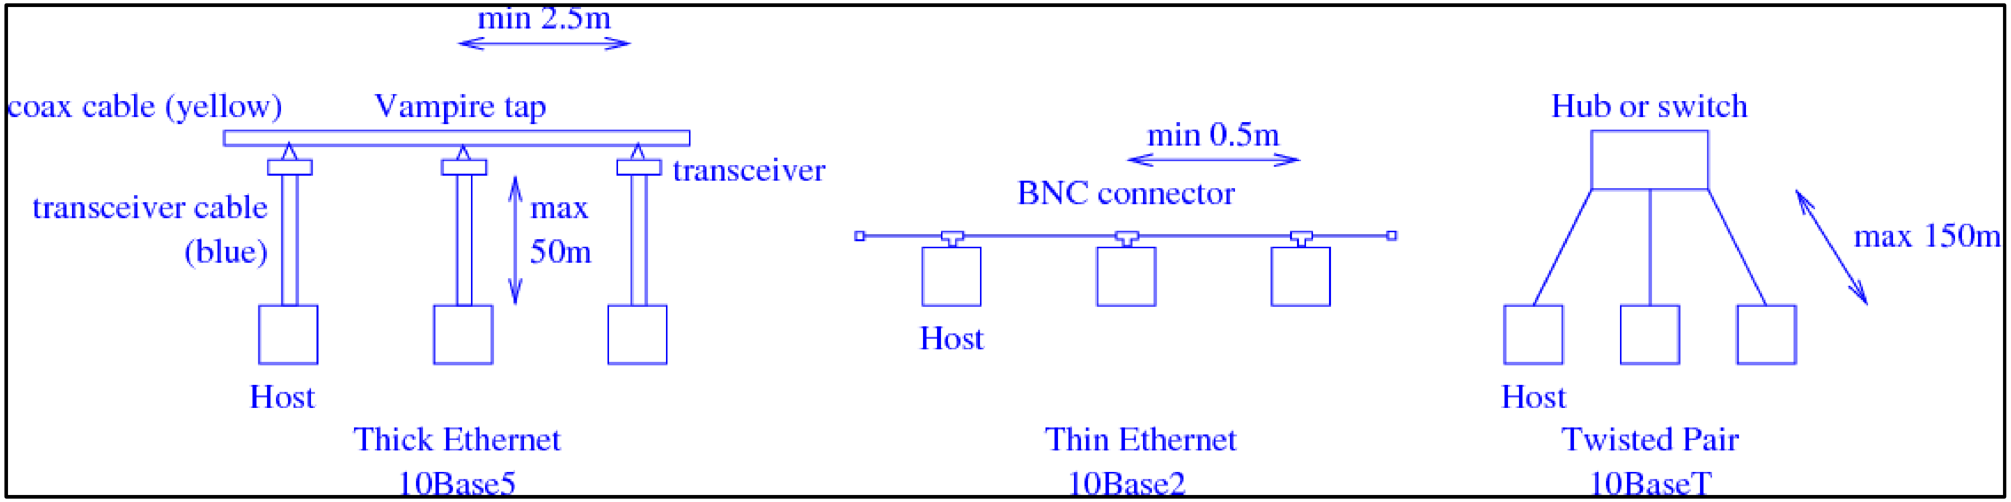

10Base5	Thick coax	500m	10Mb/s
10Base2	Thin coax	200m	10Mb/s
10BaseT	Twisted pair	100m	10Mb/s
10BaseF	Fibre optic	2000m	10Mb/s

Ethernet Hardware

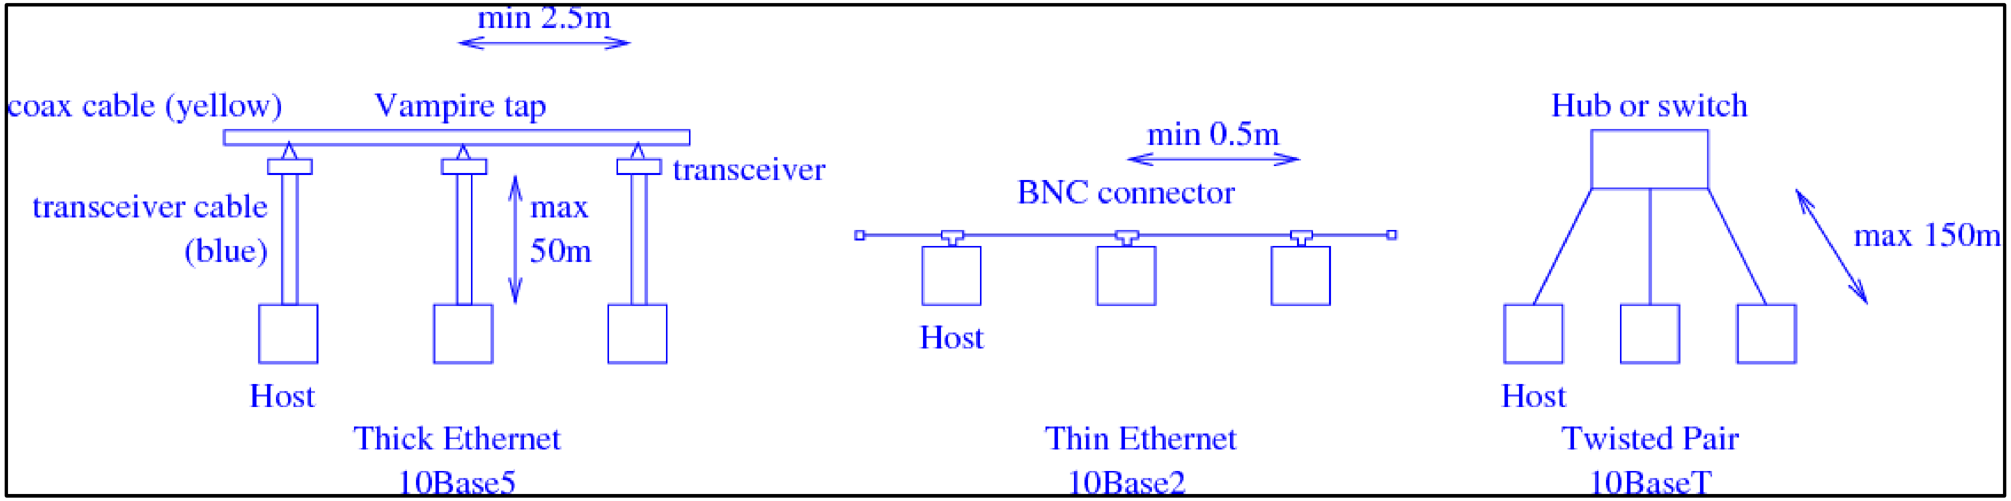

100BaseT4	Twisted pair	100m	100Mb/s
100BaseT	Twisted pair	100m	100Mb/s
100BaseF	Fibre optic	2000m	100Mb/s
1000BaseT	Twisted pair	100m	1Gb/s

- Plus several others

Ethernet Hardware

- 10Base5: fat coaxial cable (yellow) with *vampire taps* and *drop cables* (blue) to the hosts, AUI plugs
- 10Base2: Thinnet, simple coaxial cable, BNC connectors directly to the hosts
- 10BaseT: twisted pair (UTP), RJ45 plugs, each host connects to a central *hub*

Ethernet Hardware

Unshielded Twisted Pair (UTP)

- Category 1: No performance criteria
- Category 2: Rated to 1 MHz (used for telephone wiring)
- Category 3: Rated to 16 MHz (used for Ethernet 10BaseT)
- Category 4: Rated to 20 MHz (used for Token-Ring, 10BaseT)
- Category 5: Rated to 100 MHz (used for 1000BaseT, 100BaseT, 10BaseT)

Ethernet Hardware

Unshielded Twisted Pair (UTP)

- Enhanced Category 5: Rated to 200 MHz (used for 1000BaseT, 100BaseT, 10BaseT)
- Category 6: Rated to 250 MHz (used for 1000BaseT)

Ethernet Hardware

Unshielded Twisted Pair (UTP)

- Enhanced Category 5: Rated to 200 MHz (used for 1000BaseT, 100BaseT, 10BaseT)
- Category 6: Rated to 250 MHz (used for 1000BaseT)
- Category 7: for the future, but should be *shielded* twisted pairs (STP), new connectors (GG45), 600MHz

Ethernet Hardware

Hubs and Switches

- A *hub* simply echoes all inputs to all outputs
- Provides a single *collision domain*
- The available bandwidth shared between all the hosts

Ethernet Hardware

Hubs and Switches

- A *switch* understands the link layer and forwards a packet to the appropriate single output
- Each output cable is now a separate collision domain
- The full bandwidth available on *each* output
- Collisions only if two hosts send to the same destination simultaneously

Ethernet Hardware

Hubs and Switches

- Switches can *store and forward* packets
- Then there can be *no* collisions and we can do away with CSMA/CD
- Buffers can fill up, though, then packets will be dropped

Ethernet Hardware

Hubs and Switches

- Switches can *cut through*, sending the start of the packet onwards before the tail has arrived
- Less latency through the switch, but would forward corrupted packets

Ethernet Hardware

Hubs and Switches

- Switches can run *full duplex*, with independent inward and outward traffic to each host
- This gives twice the total bandwidth
- No collisions possible as inward and outward traffic runs over different twisted pairs

Ethernet Hardware

Further and Faster

- 10Mb/s, 100Mb/s, 1Gb/s, ...
- These can *autonegotiate* to select optimum speed
- Gigabit over copper: very complicated hardware
- Gigabit CSMA/CD requires *carrier extension* to make the packets big enough
- Compensates with *packet bursting*

Ethernet Hardware

Further and Faster

- 10Gb/s Ethernet coming soon
- Full duplex switched only, no CSMA/CD
- Mainly fibre optic, but copper is under consideration
- 40Gb/s and 100Gb/s Ethernet in the planning stages
- Proponents claim that Ethernet will take over the world!

Ethernet Hardware

Physical Encodings

- 0V for 0 and 1V for 1?
- Empty network and stream of 0s looks the same
- Bits need to be synchronised to prevent drifting out of step (was that 1000 or 999 0s?)
- A long stream of 1s is a steady 1V: this is electrically a bad design, an average 0V is best

Ethernet Hardware

Manchester Encoding

- Split time interval for a bit into two parts
  - Low then high voltage is a 0
  - High then low voltage is a 1
- So average is 0V
- -0.85V for low, +0.85V for high
- Easy to synchronise: transit through 0V is the middle of a bit
- But doubles the frequency of the signal to 20Mhz

Ethernet Hardware

Manchester Encoding?

- What of 100Mb/s Ethernet?
- Can't use Cat 5 cables with Manchester as only specified to 100MHz, need 200Hz

Ethernet Hardware

4B/5B Encoding

- Instead encode 4 data bits as 5 physical bits
- E.g., 0000 become 11110

Input	4B/5B	Input	4B/5B
0000	11110	1000	10010
0001	01001	1001	10011
0010	10100	1010	10110
0011	10101	1011	10111
0100	01010	1100	11010
0101	01011	1101	11011
0110	01110	1110	11100
0111	01111	1111	11101

- IDLE 11111; QUIET 00000; HALT 00100; etc.

Ethernet Hardware

Symbols

- A physical representation is called a *symbol*
- Need not be binary
- Need not represent a whole number of bits
- The *baud rate* is the number of symbols per second

Ethernet Hardware

4B/5B Encoding

- Instead encode 4 data bits as 5 physical bits

E.g., 0000 become 11110

- This has made things worse?
- Use a *three* level physical encoding *MLT-3*
- This has +, 0, and – levels (±0.85V), again using transitions to encode bits

Ethernet Hardware

MLT-3 Encoding

- Transitions are cyclically – to 0, 0 to +, + to 0, and 0 to –
- A transition marks a 1, no transition marks a 0
- The 4B/5B translation ensures that every chunk of 5 symbols has at least two transitions, so average voltage is about 0
- E.g., input 0000, with no transitions becomes 11110 with four transitions

Ethernet Hardware

MLT-3 Encoding

- Example (hex) value 0E = 0000 1110

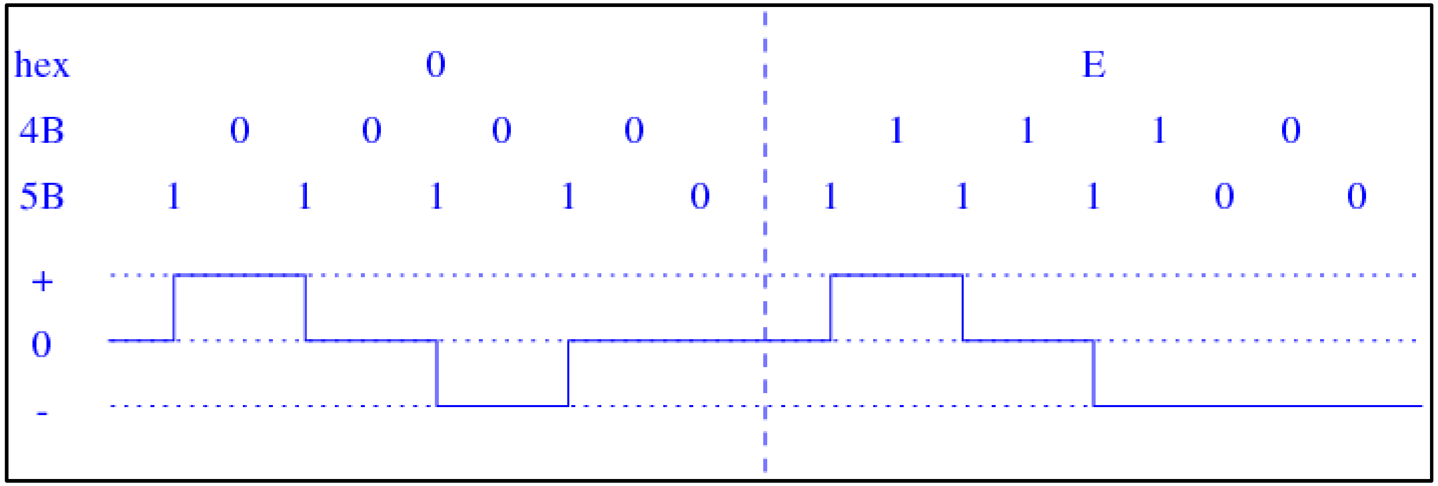

Ethernet Hardware

MLT-3 Encoding

- This runs at up to 31.25MHz for a symbol rate of 125MBaud: all 1s output (IDLE) is four transitions (- to 0, 0 to +, + to 0, 0 to -) per cycle
- Data symbols have four 1s, giving a frequency of 25MHz
- This has a symbol rate of 125MBaud for a data rate of 100Mb/s: 80% efficient or 1 symbol is 0.8 bits

Ethernet Hardware

Gigabit Encoding

- 8 data bits become 10 physical bits
- A 5 level encoding (PAM-5), ±2, ±1, 0 encodes 2 bits per symbol (plus one symbol spare)
- Over all four pairs in the cable simultaneously (in both directions on all pairs)
- 125MBaud rate, so 2 bits x 125MBaud x 4 pairs = 1000Mb/s

Token Ring

An alternative to Ethernet

- A collection of point-to-point links that form a ring

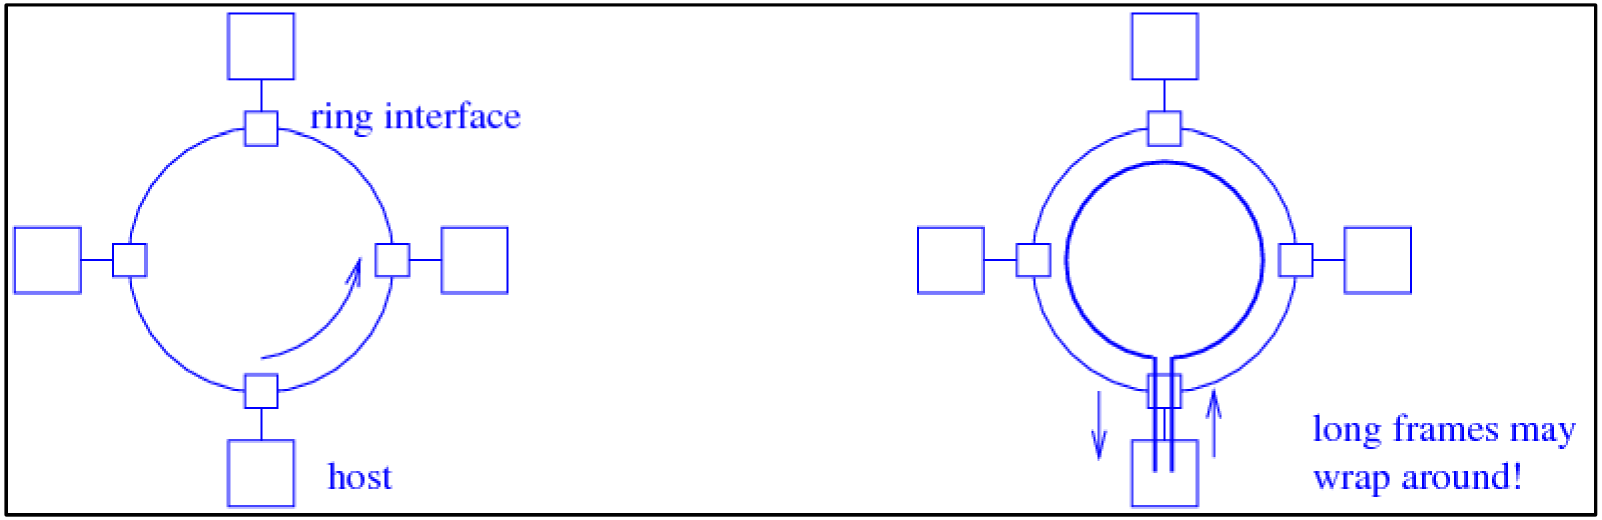

Token Ring

- A packet circulates around the ring from source to destination, passing from one interface to the next
- Ring is controlled by a special packet, the *token*
- When the ring is idle, the token circulates
- When a host wants to send it must wait for the token to arrive
- It removes the token and replaces it by the data packet

Token Ring

- The data packet is passed to the destination, which reads it and marks it as read
- The data packet *carries on back to the source*
- The source sees the packet, notes it has been read successfully, and replaces it with the token
- The token implements a natural sharing of the network in a round-robin style

Token Ring

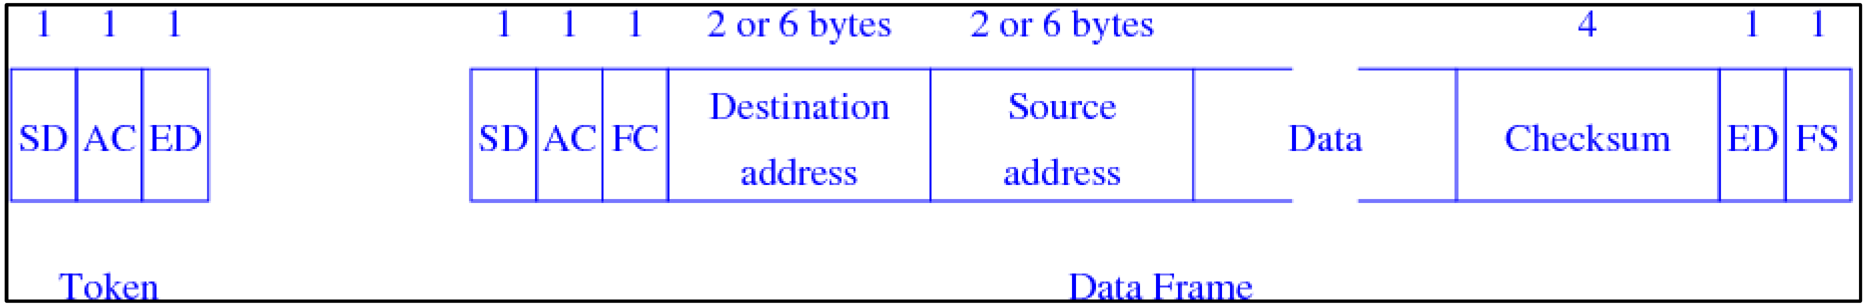

- A packet is of arbitrary length
- Starting delimiter; access control; ending delimiter
- Frame control; addresses; the data
- Checksum; ending delimiter; frame status

Token Ring vs Ethernet

- Token ring has good bounds for time to access network; Ethernet is non-deterministic
- Token ring degrades smoothly with load; Ethernet gets very inefficient
- Token rings can get very large; Ethernets are limited in size
- Token rings have no minimum packet size

Token Ring vs Ethernet

- Token rings are difficult to extend: the entire ring must stop to add a new node
- A broken interface in a Token ring stops the whole ring
- There is a large possible delay while waiting for the token; Ethernet is near instant at low load
- Token rings require complex management (regenerating lost tokens, etc.); Ethernet is easy

Token Ring vs Ethernet

- Ethernet is everywhere; Token is ring rarely seen