Ethernet (10 Mbps shared)

Adam M. Costello <amc@cs.berkeley.edu> 1999-Sep-25-Sat

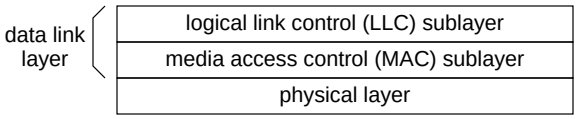

Hardware:

Originally coaxial cable connected to multiple network interfaces. Repeaters can make multiple cables look like one (every interface can hear every other interface).
Now interfaces are usually connected by unshielded twisted pairs of copper wires directly to hubs (repeaters with many ports).

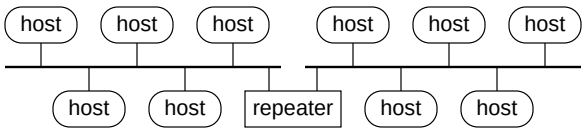

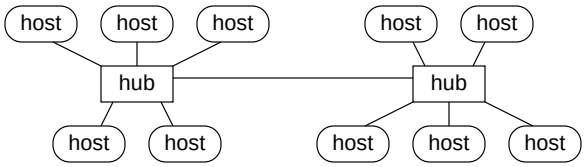

Transmitting bits:

Must transmit both the bits and a clock signal, so receiver knows when to sample the bits. Manchester encoding puts a transition in the middle of each bit:

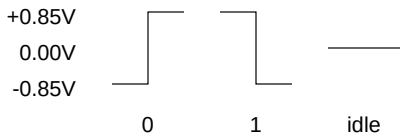

Receiver uses the transitions to keep its clock in sync with the data. A frame is preceeded by a preamble (62 bits of 101010...) and then a start-of-frame delimiter (11). The end of the frame is marked by a return to idle.

MAC frame:

destination address	6 bytes
source address	6 bytes
type/length	2 bytes
data	46 to 1500 bytes
CRC (checksum)	4 bytes

If type/length is more than 1500, this is a DIX frame (an Ethernet frame as originally defined by DEC, Intel, and Xerox). The type indicates what kind of payload the data is (like IP datagram, ARP message, AppleTalk packet, etc). Notice that the end of a small payload padded to 46 bytes is not indicated; such payloads must be self-delimiting.

If type/length is 1500 or less, this is an IEEE 802.3 frame. The data is an LLC frame, and the length field is its length in bytes.

Carrier sense multiple access (CSMA):

If two interfaces transmit at the same time, both signals will be garbled. To avoid stepping on each other's frames, interfaces wait for the wire to be idle before beginning to transmit.

Collision detection (CD):

Two interfaces might start transmitting at about the same time, causing a collision. An interface monitors the line while transmitting. If it hears a collision, it sends a jam signal (32 random bits) to make sure the other transmitters hear the collision, then stops transmitting and tries again a little later.

We want a transmitter to hear a collision before it stops transmitting, so it knows to try again. So the frame duration must be at least the maximum round-trip time. Hence the minimum frame size, and limits on cable length and number of chained repeaters.

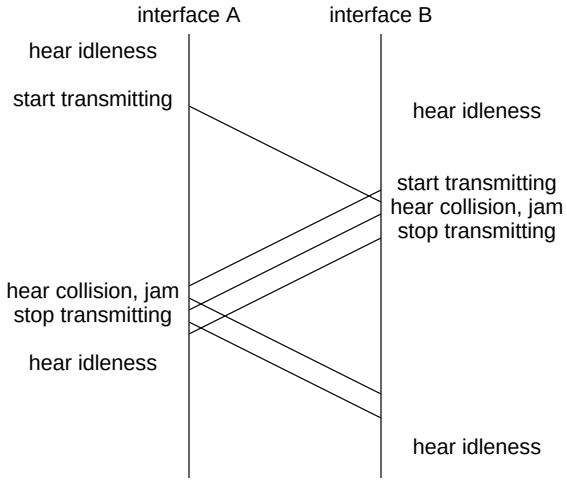

Exponential backoff:

After the nth collision with the same frame, wait K slot-times before trying again, where K is uniform random in [0, (2^n)-1], or [0,1023] if n > 10 because no more than 1024 hosts are allowed on a shared Ethernet. A slot time is 51.2 microseconds (512 bit-times), which is longer than the maximum round-trip time, so if the smallest K picked by the contenders is unique, everyone will hear that transmitter and not collide with it. After 16 tries the interface gives up on that frame.

A wide interval is fast for many contenders, slow for few contenders. A narrow interval is fast for few contenders, slow for many. The backoff algorithm adapts to the current number of contenders.

LLC frame:

destination service access point (DSAP)	1 byte
source service access point (SSAP)	1 byte
control	1 byte

DSAP and SSAP identify the payload type.
The LLC sublayer can optionally provide reliability, but it is not used on the Internet where packet losses on wired LANs are negligible compared to losses due to congestion in routers.