CS 123: Lecture 12, LANs, and Ethernet

George Varghese

October 24, 2006

Selective Reject Modulus failure Example

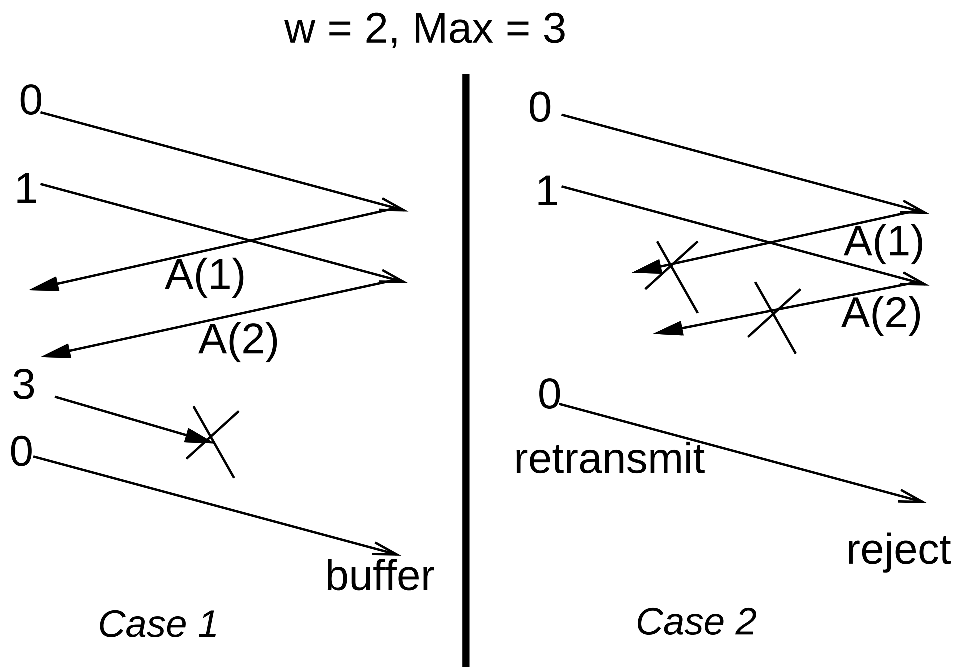

Design Lessons

- Prefer correctness not to depend on timer settings.
- Don't add complexity till its justified. Stop and wait to sel reject.
- Design Simple Protocols first and optimize later.
- Consider dynamic versions of static protocols. Dynamic windows.
- Need to design protocols to recover from faults. Use restart method.
- Impossibility Results teach us what we have to change to do our jobs.

Course Strategy

- **Build concepts systematically:** layer upon layer, sublayer upon sublayer starting from bottom.
- **Learn some useful, general ideas:** Clock recovery, synchronization, sublayering, coding, Shannon Limits, Hamming Distance, latency versus throughput, state view and invariants.
- **Question existing mechanisms:** Why CRC-16, why mod 2 division, why is HDLC flag 01111110?
- **Learn to Extend and Create:** Creating stuffing schemes, new error detection schemes, new error recovery schemes, new transports and routing schemes later.

A Rose by any other Name

Inherently sequential. Link shared among multiple senders. A.K.A.

- *Multiaccess links* because there are multiple nodes that may simultaneously access link.
- *Broadcast links* because every transmission can be heard by all other stations.
- *Local Area Networks* or LANs because the geographical area serviced by a LAN is local and small – from 1-10 Kms, covering an office, a building, or at most a campus.

WHY LANS

- Connect up all computers in a building, office, campus. 200-1000 users. Saves wiring costs to share one wire.
- Provides high bandwidth and low error rates for local group of users. Worth it because most high-bandwidth distributed computing (e.g., file servers etc.) has acess locality.
- Time division multiplexing not a good idea when user traffic is bursty as data is. Bursty = high peak/average ratio. Each user gets access to entire LAN bandwidth when other users are idle. As more users are added, they share the bandwidth

Statistical versus Strict Multiplexing

- Strict multiplexing: TDM or FDM where a user is given a fixed allocation *regardless* of whether the user has data to send or not.
- *Bursty*: traffic has a high peak/average ratio.
- Strict gives each user B/N, where N is the number of possible sources; stat multiplexing attempts to give each user, B/x, where x is the number of busy users. N large (100 -1000) while x is small (1-10).
- Suppose 100 users each transfer a 125 kbyte file every 2 hours. At 10 Mbps, a 125 kbyte file takes 0.1 seconds to transmit. At 0.1 Mbps, file takes 10 seconds to transfer.

ALOHA

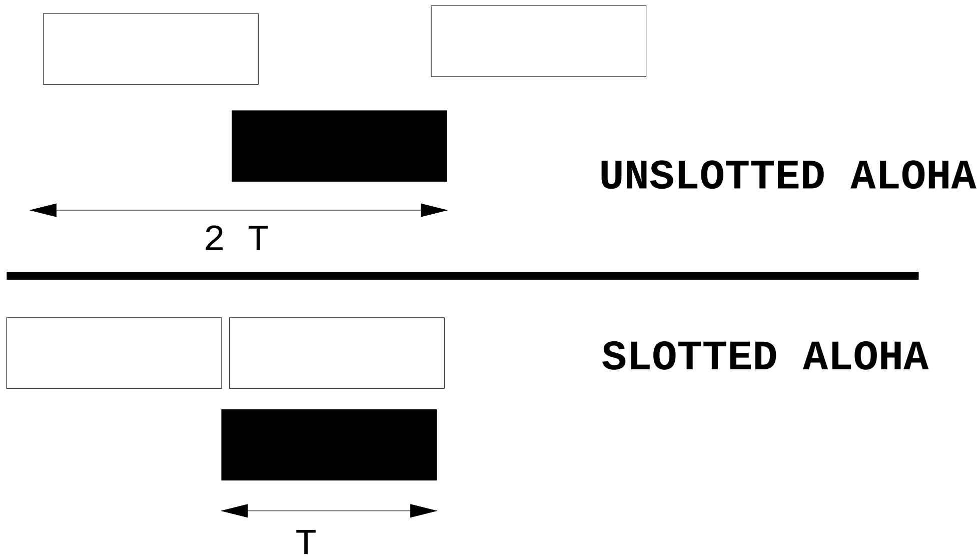

- Predecessor of Ethernet. Multiple ground stations in various parts of Hawai.
- Couldn't do collision detection or carrier sense, but did have a form of randomized backoff. More inefficient: but similar problems arise in mobile computing.
- Slotted Aloha reduces vulnerable period by half but requires a common clock reference.

In transmission Collision Detect and Semi-reliability

- 1500 byte frame involved in a collision. Ethernet *aborts* transmission after 64 bytes. Aloha will send the entire 1500 bytes and detect when ack is not received. Better for large frames, and large frame sizes allow efficiency.
- No detection of frame corruption. (1 in million). However, collisions are frequent. Semi-reliable: detect collisions and retransmit in hardware.

ETHERNET

Three mechanisms: :

- Carrier sense and deference: No point transmitting when someone else is speaking
- Collision Detection: No point continuing to send a large frame (up to 1500 bytes) when after 64 bytes you can detect a collision.
- Exponential Backoff: Collisions very frequent, so important to retransmit. Random backoff avoids synchronized collisions. Exponentially growing backoff attempts to dynamically adjust to number of colliders. Fixed backoff would be wasteful or insufficient.

Other details:

- Slot time = 2T, T = end-to-end delay = 51.2 usec = 64 bit times = maximum delay to detect a collision.
- Minimum packet size of 64 bytes to avoid possibility of finishing transmission before collision is detected. Add pad if data is too small.
- Jam: transmit small number of bits after you detect a collision to ensure that other transmitters also detect collision.
- Use Manchester with average DC level per bit. Collision detection by detecting increased voltage level.

Dynamic Backoff

- Consider 2 colliders. One should wait 1 slot and the other 0. Can approximate by tossing a coin.
- Consider 16 colliders. Ordinary coin-tossing does not work. Pick random numbers from 1 to 16.
- static scheme either does not work for a large number of colliders or is inefficient for a small number.
- Binary exponential backoff.

Why Min Packet Sizes

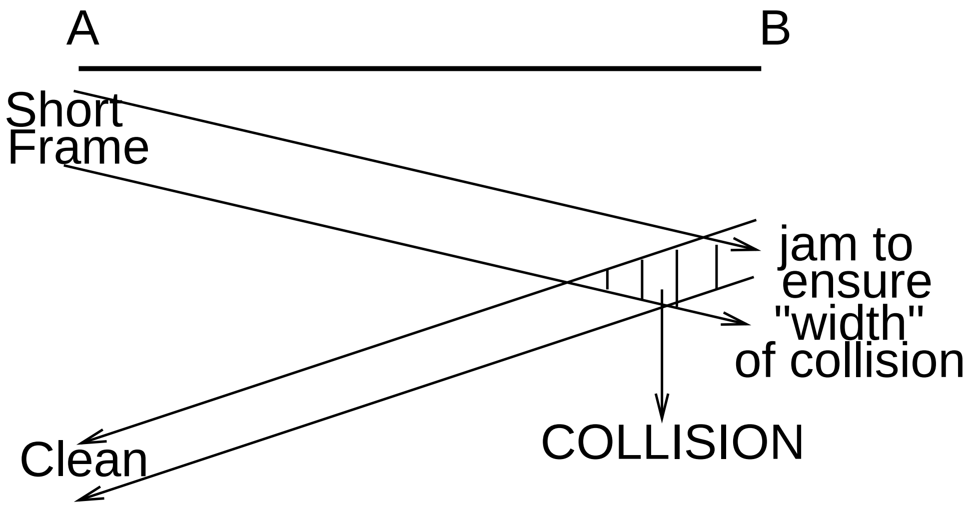

•

•

Ethernet Header

01010111 preamble	Dest (6)	Source (6)	Length (2)	Data	Pad	CRC

Total Frame length 64<=L<=1500

LLC sublayer

Ethernet Implementation Details

- Limited distance (2.5 km). 500 m wires. 4 repeaters.
- Thin wire, Thick wire
- Repeaters: important device, reads in a bit and writes out a bit on other side, boosting signal strength.
- Physical Topology is a star or tree.

Cost of statistical muxing in Ethernet

- If distance goes up by a factor of 10 or speed, then what happens to min packet size? Why is it wasteful?
- New 100 Mbps Ethernet proposal, has only 200 m extent.

Concepts

- Statistical Multiplexing
- Importance of Pipe size in determining efficiency.
- Logical versus physical topology