CS244a: An Introduction to Computer Networks

Handout 10: Link Layer
CSMA/CD, Ethernet, Token Passing

Nick McKeown
Professor of Electrical Engineering and Computer Science, Stanford University
nickm@stanford.edu
http://www.stanford.edu/~nickm

The Link Layer

Application
Presentation
Session
Transport
Network
Link
Physical

The 7-layer OSI Model

Telnet FTP SMTP HTTP NNTP	TFTP
TCP	UDP
IP	
LAN-LINK	

The 4-layer Internet Model

Examples of MAC Protocols

(MAC = "Medium Access Control")

Simple
Random

Packet-Switched Radio Network

Aloha

Carrier Sense Multiple Access/Collision Detection

Ethernet (IEEE 802.3)

Complex
Deterministic

Token Passing

Token Ring (IEEE 802.5)

Goals of MAC Protocols

MAC Protocols arbitrate access to a common shared channel among a population of users

1. *Fair among users*
2. *High efficiency*
3. *Low delay*
4. *Fault tolerant*

Outline

- Random Protocols
  - Aloha
  - CSMA/CD
  - Ethernet (CSMA/CD put into practice)
- Token Passing Protocols
  - Common Features
  - Flavor #1: Release After Reception (RAR)
  - Flavor #2: Release After Transmission (RAT)

Aloha Protocol

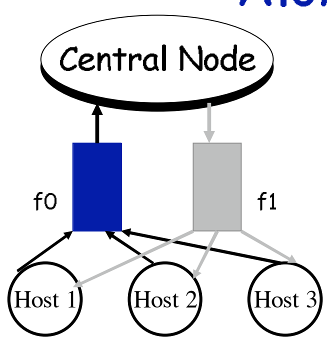

Basic operation:
1. All hosts transmit on one frequency.
2. Central node *repeats* whatever it receives on the other frequency.

If more than one host transmits at the same time
➡ ***Collision at central node!***

If there is a collision, hosts receive corrupted data, and so wait for a randomly chosen time before retransmitting their packets.

Aloha Protocol

- Aloha protocol is very simple, and fairly robust against failure of a host.
- The protocol is distributed among the hosts.
- Under low-load, we can expect the delay to be small.
- Under high-load, a lot of time is "wasted" sending packets that collide.

Improving performance:

- Listen for activity before sending a packet.
- Detect collisions quickly and stop transmitting.
- After a collision, pick the random waiting time so as to maximize throughput.

CSMA/CD Protocol

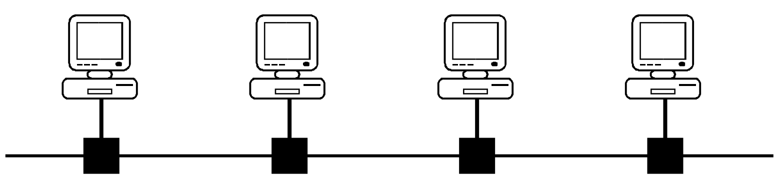

All hosts transmit & receive on one channel
Packets are of variable size.

When a host has a packet to transmit:

1. *Carrier Sense*: Check that the line is quiet before transmitting.
2. *Collision Detection*: Detect collision as soon as possible. If a collision is detected, stop transmitting; wait a *random time*, then return to step 1.

binary exponential backoff (↑ random time)

CSMA/CD Network Size Restriction

To ensure that a packet is transmitted without a collision, a host must be able to detect a collision before it finishes transmitting a packet.

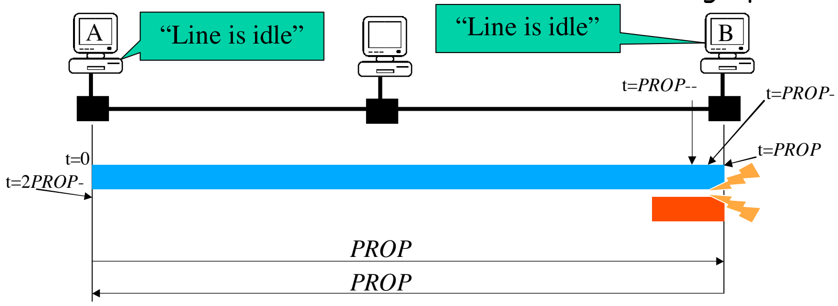

Events:

t=0: Host A starts transmitting a packet.

t=*PROP*--: Just before the first bit reaches Host B, Host B senses the line to be idle and starts to transmit a packet.

t=*PROP*-: A collision takes place near Host B.

t=*PROP*: Host B receives data whilst transmitting, and so detects the collision.

t=2*PROP*-: Host A receives data whilst transmitting, and so detects the collision.

CSMA/CD Network Size Restriction

"To ensure that a packet is transmitted without a collision, a host must be able to detect a collision before it finishes transmitting a packet."

From example on previous slide we can see that for a Host to detect a collision before it finishes transmitting a packet, we require:

$$TRANSP > 2 \times PROP$$

In other words, there is a minimum length packet for CSMA/CD networks.

Performance of CSMA/CD

We're going to analyze the performance of a CSMA/ CD network.

- Our performance metric will be Efficiency, η. This is defined to be the fraction of time spent sending useful/successful data. The more time spent causing and detecting collisions, the less efficient the protocol is. More precisely:

$$\eta = \frac{\text{Time taken to send data}}{\text{Time taken to send data + overhead}}$$

- To make the analysis simple, we'll assume that time is slotted and all packets are the same length. A time slot equals 2 x PROP. In any given time slot, a host will either decide to transmit or not with probability p. (This includes packets transmitted for the first time and retransmissions).
- First, we will try and find the value of p that maximizes the throughput (in fact, it's the goodput).
- Then, using the optimal value of p, we'll find the efficiency.

Performance of CSMA/CD

Maximizing goodput

Find the goodput, $\alpha(p)$:

Probability that exactly one node transmits in a given slot.

$$\alpha(p) \equiv \binom{N}{1} p(1-p)^{N-1}$$

$$\frac{d\alpha}{dp} = N(1-p)^{N-1} - pN(N-1)(1-p)^{N-2}$$

$$\therefore \alpha_{\max} \approx 36\% \approx 40\% \qquad when: \qquad p = 1/N$$

Performance of CSMA/CD

Finding the overhead

Define A to be the expected number of time slots wasted before a packet is transmitted successfully:

$$A = (\alpha \times 0) + (1-\alpha)(1+A)$$

$$\therefore when: \qquad \alpha = \alpha_{\max}, \qquad A = 1.5$$

[Alternatively, consider a coin with Pr(heads) = α = 0.4. The expected number of coin tosses until the first head is 1/0.4 = 2.5.
i.e. 1.5 unsuccessful attempts, followed by 1 successful one]

Performance of CSMA/CD

Finding the efficiency

$$\eta_{CSMA/CD} = \frac{TRANSP}{TRANSP + E[\text{# of wasted slots per packet}]}$$

$$= \frac{TRANSP}{TRANSP + A(2 \times PROP)}$$

$$= \frac{TRANSP}{TRANSP + (3 \times PROP)}$$

Performance of CSMA/CD

$$\eta_{CSMA/CD} = \frac{1}{1+3a}, \quad \text{where: } a \equiv \frac{PROP}{TRANSP}$$

From simulation and more precise models:

$$\eta_{CSMA/CD} \approx \frac{1}{1+5a}$$

Outline

- Random Protocols
  - Aloha
  - CSMA/CD
  - Ethernet (CSMA/CD put into practice)
- Token Passing Protocols
  - Common Features
  - Flavor #1: Release After Reception (RAR)
  - Flavor #2: Release After Transmission (RAT)

The Original Ethernet

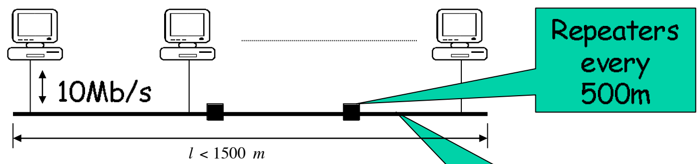

$$PROP_{max} = l / c = 1500 / 2.5 \times 10^{8} = 6 \mu s$$
$$TRANSP > 2PROP \Rightarrow TRANSP > 12 \mu s$$
$$\therefore Packetsize \geq (12 \mu s) \times 10 Mb/s = 120 bits$$

In practice, minimum packet size = 512 bits.

- allows for extra time to detect collisions.
- allows for "repeaters" that can boost signal.

The Original Ethernet

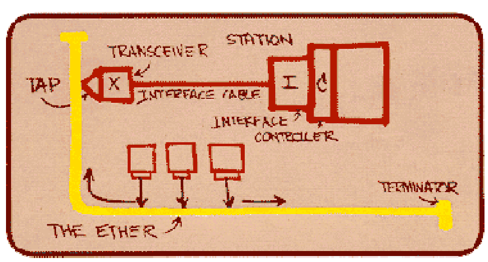

Original picture drawn by Bob Metcalfe, inventor of Ethernet (1972 – Xerox PARC)

Ethernet Frame Format

Bytes:	7	1	6	6	2	0-1500	0-46	4
	Preamble	SFD	DA	SA	Type	Data	Pad	CRC

1. Preamble: trains clock-recovery circuits
2. Start of Frame Delimiter: indicates start of frame
3. Destination Address: 48-bit globally unique address assigned by manufacturer.
   - 1b: unicast/multicast
   - 1b: local/global address
4. Type: Indicates protocol of encapsulated data (e.g. IP = 0x0800)
5. Pad: Zeroes used to ensure minimum frame length
6. Cyclic Redundancy Check: check sequence to detect bit errors.

The 10Mb/s Ethernet Standard

IEEE 802.3

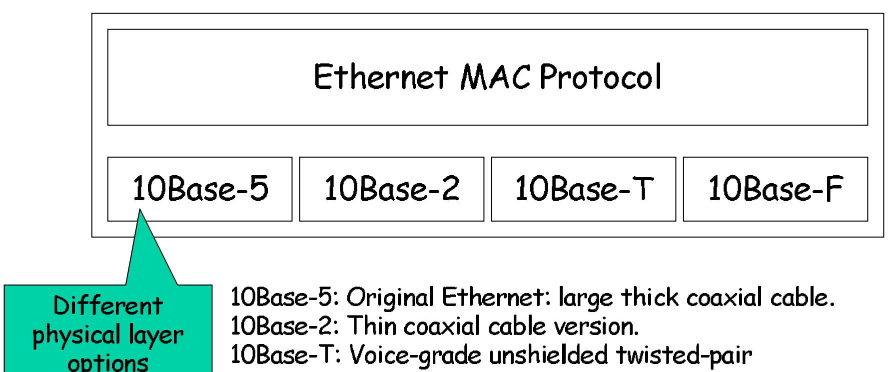

10Base-5: Original Ethernet: large thick coaxial cable.
10Base-2: Thin coaxial cable version.
10Base-T: Voice-grade unshielded twisted-pair Category-3 telephone cable.
10Base-F: Two optical fibers in a single cable.

10Base-T
"Twisted pair Ethernet"

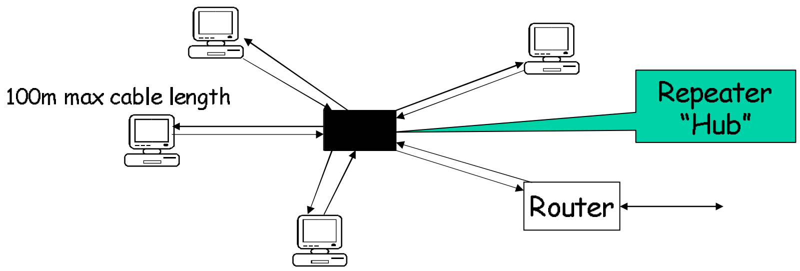

- Designed to run over existing voice-grade "Category-3" twisted pair telephone wire.
- Centralized management ("managed hubs") lead to more reliability.
- Created a huge increase in Ethernet usage.

Increasing the data rate
10Mb/s -> 100Mb/s -> 1Gb/s -> 10Gb/s

- Problem: $TRANSP > 2PROP$
- E.g. CSMA/CD at 100Mb/s over 1500m of cable: $PROP = 1500 / 2.5 \times 10^8 = 6\mu s$

  $\therefore TRANSP > 12\mu s \Rightarrow Packetsize \geq 1200bits$
- To overcome this two techniques used:
  - Cable length limited to 100m:

    $PROP = 200 / 2.5 \times 10^8 \Rightarrow Packetsize \geq 160bits$
  - Use "Ethernet Switching" to prevent collisions (in an upcoming lecture).

The 100Mb/s Ethernet Standard

"Fast Ethernet"

Different physical layer options

Ethernet MAC Protocol

100Base-T4 | 100Base-TX | 100Base-FX

Up to 100m of cable per segment.

100Base-T4: Uses four pairs of voice grade Category-3 cable .
100Base-TX: Uses two pairs of data grade Category-5 cable.
100Base-FX: Uses two optical fibers.

The 1Gb/s Ethernet Standard

"Gigabit Ethernet"

Ethernet MAC Protocol

1000Base-TX | 1000Base-FX

1000Base-TX: Uses four pairs of data grade Category-5 cable.
1000Base-FX: Uses two optical fibers.

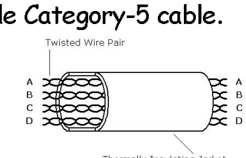

Outline

- Random Protocols
  - Aloha
  - CSMA/CD
  - Ethernet (CSMA/CD put into practice)
- Token Passing Protocols
  - Common Features
  - Flavor #1: Release After Reception (RAR)
  - Flavor #2: Release After Transmission (RAT)

Token Passing
Common Features

- A token rotates around a ring to each node in turn. We will define:
  PROP = minimum rotation time around ring.
- All nodes (computers, routers, etc.) *copy* all data and tokens, and *repeat* them along the ring.
- When a node wishes to transmit packet(s), it *grabs* the token as it passes.
- It *holds* the token while it transmits.
- When it is done, it *releases* the token again and sends it on its way.

Token Passing

Common Features

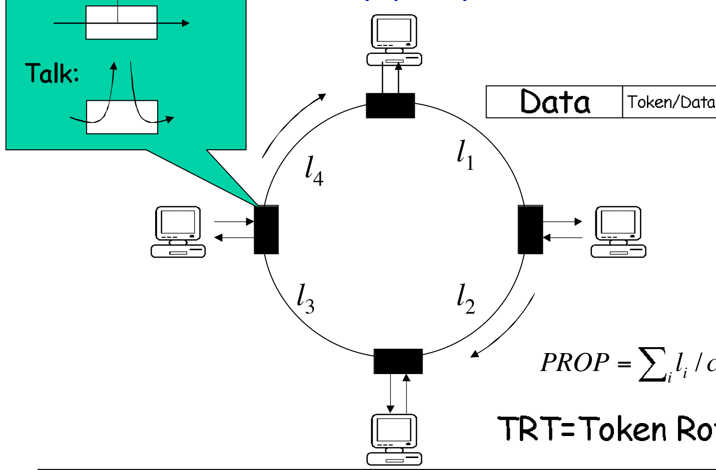

$$PROP = \sum_i l_i / c = TRT_{min}$$

TRT=Token Rotation Time

Flavor #1: Release After Reception (RAR)

- Computer captures token, transmits data, waits for data to successfully travel around ring, then releases token again.
- Allows computer to detect errored frames and retransmit them.

Example time evolution in which host 1 and host 3 have packets to transmit:

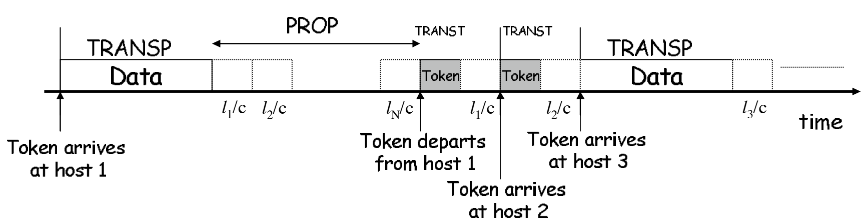

Efficiency of RAR

Recall: Efficiency, η, is the fraction of time spent sending useful data.

Define: $T_{i,j}$ to be the time from when the token arrives at host i until it next arrives at host j.

$$T_{1,2} \le TRANSP + PROP + TRANST + l_1 / c$$
$$\therefore T_{1,1} \le N(TRANSP + PROP + TRANST) + \sum_i l_i / c$$
$$= N(TRANSP + PROP + TRANST) + PROP$$

$$\therefore \eta_{RAR} \le \frac{N(TRANSP)}{N(TRANSP + PROP + TRANST) + PROP}$$
$$\approx \frac{1}{1+a}, \quad a = \frac{PROP}{TRANSP}, \quad TRANSP >> TRANST$$

Flavor #2: Release After Transmission (RAT)

- Computer captures token, transmits data, then releases token again.

Example time evolution in which host 1 and host 3 have packets to transmit:

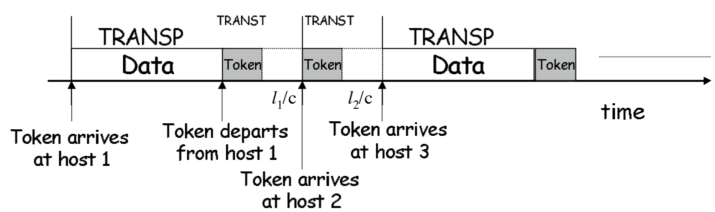

Efficiency of RAT

$$T_{1,2} \le TRANSP + TRANST + l_1 / c$$
$$\therefore T_{1,1} \le N(TRANSP + TRANST) + \sum_i l_i / c$$
$$= N(TRANSP + TRANST) + PROP$$

$$\therefore \eta_{RAT} \le \frac{N(TRANSP)}{N(TRANSP + TRANST) + PROP}$$
$$\approx \frac{1}{1 + a / N}, \quad a = \frac{PROP}{TRANSP}, \quad TRANSP >> TRANST$$

Comparison of Efficiencies

$$\eta_{CSMA/CD} \approx \frac{1}{1+5a}$$
$$\eta_{RAR} \le \frac{1}{1+a}$$
$$\eta_{RAT} \le \frac{1}{1+a/N}$$

Example: 100 node network
- PROP = 1000m/c
- TRANSP = (1000bits)/ (100Mb/s)

$$a = 4/10 = 0.4$$
$$\eta_{CSMA/CD} \approx \frac{1}{1+2} = 33.3\%$$
$$\eta_{RAR} \le \frac{1}{1+a} = 71\%$$
$$\eta_{RAT} \le \frac{1}{1+a/N} = 99.6\%$$

Token Rings

Techniques:

- Release After Reception (RAR)
- Release After Transmissions (RAT)

Examples:

- RAR: IEEE 802.5 Token Rings
- RAT: Fiber Distributed Data Interface (FDDI)