CSMA/CD Throughput

$\longrightarrow$ approximate analysis in simplified setting

Assumptions:

- time is slotted

  $\rightarrow$ slot duration: 2τ

- k hosts; each host transmits with probability p at every slot

  $\rightarrow$ transmission behavior among hosts independent

  $\rightarrow$ transmission behavior across slots independent

Note:

- independence among users: typical assumption
- independence across time: strong assumption

CSMA/CD is a feedback control:

$\longrightarrow$ modify future behavior depending on present/past

That is:

$\rightarrow$ upon collision more to send in the future: $p \uparrow$

$\rightarrow$ upon backoff: $p \downarrow$

$\rightarrow$ more general: backlog

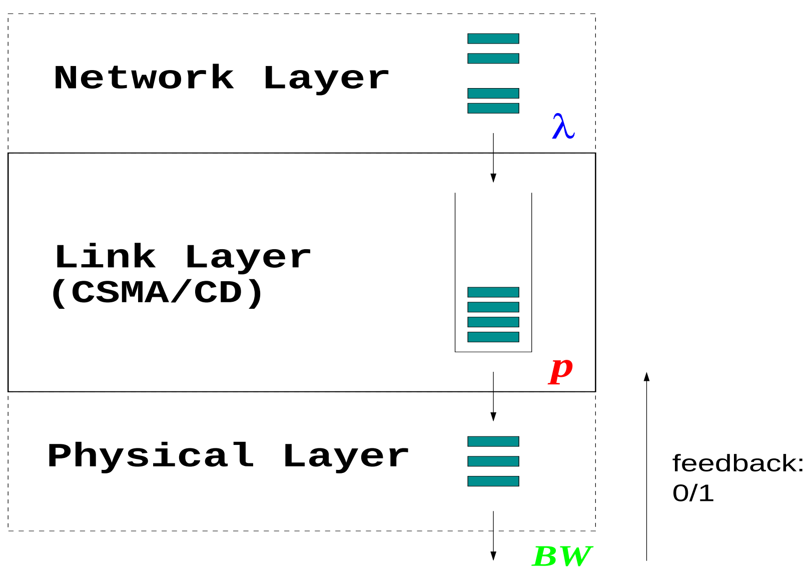

We will consider fixed, independent p

$\rightarrow$ no backlogs

$\rightarrow$ no feedback adapation of p

New performance metric: utilization (ϱ)

$\longrightarrow$ fraction of total bandwidth attained

$\longrightarrow$ $0 \leq \varrho \leq 1$

$\longrightarrow$ captures efficiency and wastage

In slotted CSMA/CD:

$\longrightarrow$ fraction of usefully used slots

$\longrightarrow$ what are "uselessly used" slots?

Ex.: snapshot of baseband channel over 10 time slots

→ blue: successfully transmitted frames

→ brown: collided frames

→ utilization ϱ?

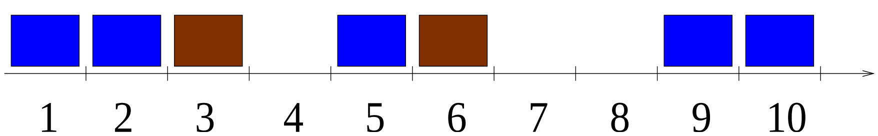

One more viewpoint:

⟶ note: useful and useless "periods" alternate

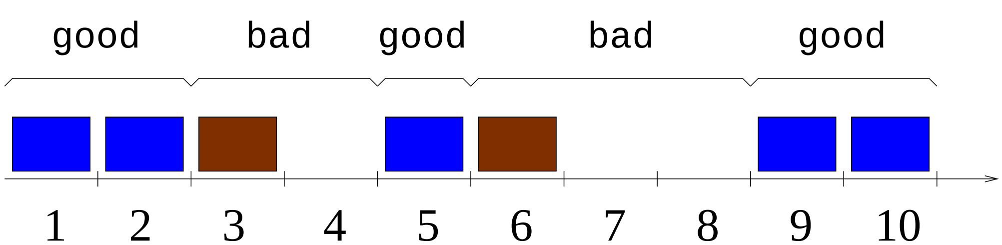

In the long run,

$$\varrho = \frac{E[\text{good}]}{E[\text{good}] + E[\text{bad}]}$$

$\longrightarrow$ avrg. length of adjacent "good" and "bad" periods

$\longrightarrow$ formula holds under mild conditions

Next: calculate $E[\text{good}]$ and $E[\text{bad}]$

Fix time slot. Probability that a fixed host acquires the slot successfully

$$p(1-p)^{k-1}$$

Probability that some host acquires the slot

$$\eta = kp(1-p)^{k-1}$$

$\longrightarrow$ why?

Now, let's be generous and find p that maximizes η

$\longrightarrow$ upper bounding

Fact: η is maximized at $p = 1/k$. Also,

$$\lim_{k\to\infty} \eta = \lim_{k\to\infty} \left(1 - \frac{1}{k}\right)^{k-1} = 1/e.$$

$\longrightarrow$ many user assumption

$\longrightarrow$ common practice to simplify expression (valid?)

Probability bad period persists for exactly i slots

$$(1-\eta)^{i-1}\eta$$

Thefore average bad period

$$E[\text{bad}] = \sum_{i=0}^{\infty} i(1-\eta)^{i-1}\eta = 1/\eta$$

$E[\text{bad}]$ is in unit of slots. Convert to second:

$$2\tau/\eta = 2\tau e$$

Similarly calculate $E[\text{good}]$; call it γ.

Convert γ to second:

$$\gamma F/B$$

where

F: frame size (bits)

B: bandwidth (bps)

Putting everything together

$$\begin{aligned} \varrho &= \frac{E[\text{good}]}{E[\text{good}] + E[\text{bad}]} \\ &= \frac{\gamma F/B}{\gamma F/B + 2\tau e} \\ &= \frac{\gamma F/B}{\gamma F/B + 2Le/c} \\ &= \frac{1}{1 + (2e/c\gamma)BL/F} \end{aligned}$$

where

L: length of wire (meters)

c: speed of light (m/s)

What does the formula say?

For example, if B is increased, what must be done to maintain high utilization?

Note: BL/F comes into play during unit conversion from slots to seconds

$\longrightarrow$ holds for different E[good] and E[bad]

$\longrightarrow$ invariant

Recall: to minimize probability of collision we set

$$p = 1/k$$

$\longrightarrow$ p-persistent version of CSMA/CD

$\longrightarrow$ depends on knowledge of number of users

$\longrightarrow$ today's Ethernet: 1-persistent

For fixed p, as a function of k CSMA/CD throughput is unimodal, i.e., dome-shaped

$\longrightarrow$ under excessive load, throughput goes down

In practice today: switched Ethernet

- contention moved from bus to "single point"

  $\rightarrow$ switch: star topology

  $\rightarrow$ analogous to old telephone switch-boards

- Ethernet frames are logically scheduled

  $\rightarrow$ includes buffering

Diagram of output-buffered switch:

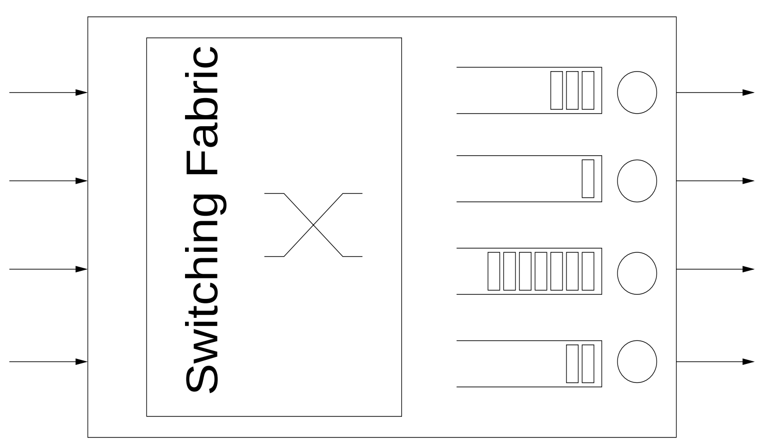

$\longrightarrow$ interconnection networks (e.g., shuffle-exchange)

$\longrightarrow$ switching fabric: hardware

- Ethernet switch emulates CSMA/CD

  → backward compatibility

  → use same frame format

- upon buffer overflow: send collision signal

  → transparent to legacy host NIC

  → awkward: instituted for incremental deployment

  → Internet: new technology must respect legacy

Ex.: 10Base-T, 100Base-T, 1000Base-T and 1000Base-X

⟶ FE: 802.3u; GigE: 802.3ab and 802.3z

⟶ negotiation: e.g., full/half duplex

⟶ how can GigE overcome length limitation?

⟶ e.g., supports 200 m as in FE

Slot time extension:

- frame format remains the same
- minimum slot time extended from 64 B to 512 B

  → padding: transparent to legacy CSMA/CA

  → also called carrier extension

  → reconciliation sublayer between MAC and PHY

Packet bursting:

- slot time extension alone problematic

  → small frames: marginal increase in throughput
- allow burst of packets

  → only first packet is padded & burst limit

Longer distances?

$\longrightarrow$ e.g., 1000Base-LX

Medium-haul GigE/10GigE (802.3ae): 500m, 5km, 40km

- CSMA/CD disabled

  → purely point-to-point link

  → switch-to-switch

  → simpler

  → backward compatibility: not an issue

- flow control

  → pause frame to prevent buffer overflow

QoS: 802.3p

$\longrightarrow$ frame tagging conveys priority

$\longrightarrow$ priority classes supported at switches

FDDI (Fiber Distributed Data Interface)

$\longrightarrow$ token ring architecture

High-bandwidth extension of IBM 4 Mbps token ring and 16 Mbps IEEE 802.5 token ring standard.

$\longrightarrow$ 100 Mbps bandwidth

Used as high-bandwidth campus/city backbone.

$\longrightarrow$ metropolitan/campus distance: MAN

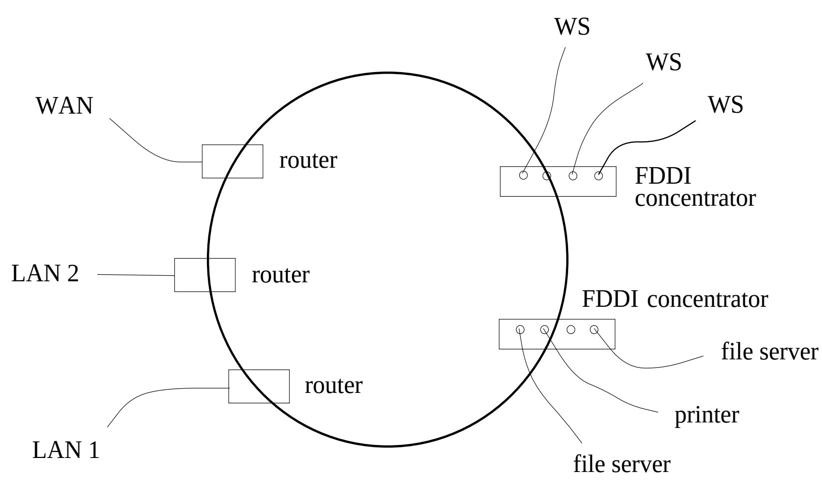

Basic operation:

$\longrightarrow$ B wants to send to D

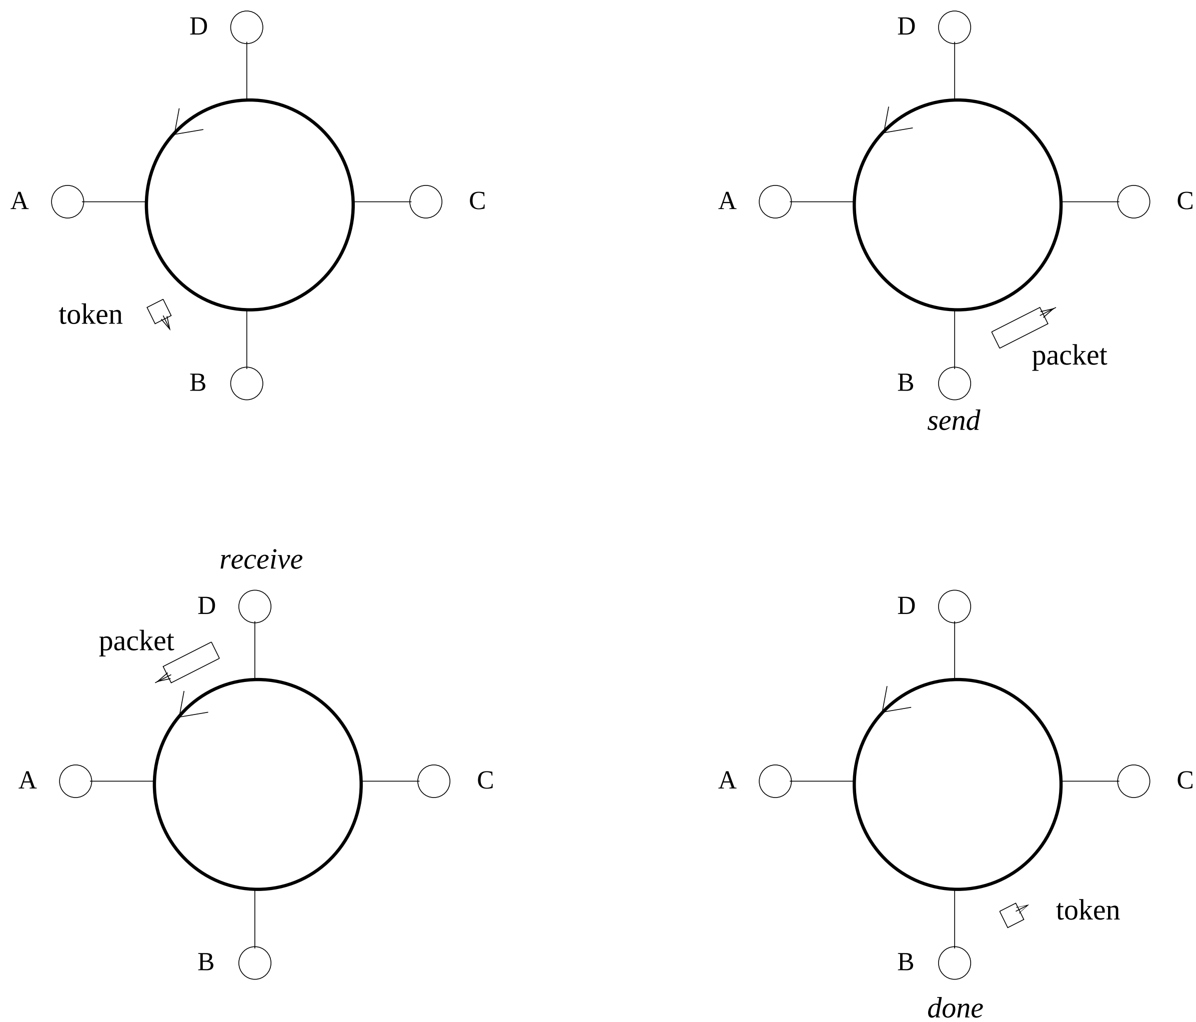

Fault-tolerance:

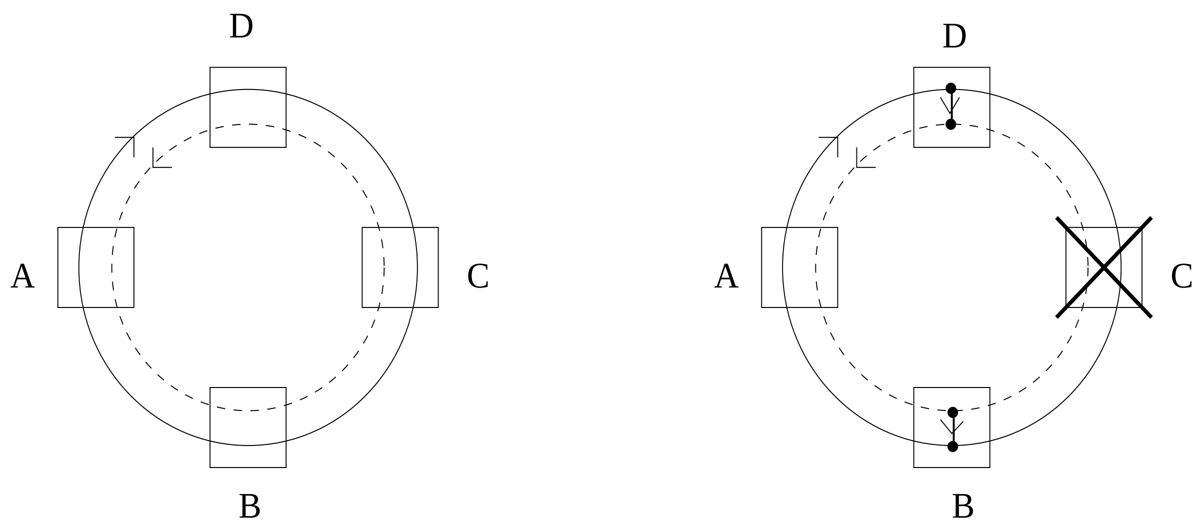

- DAS (dual attachment station)
- SAS (single attachment station)

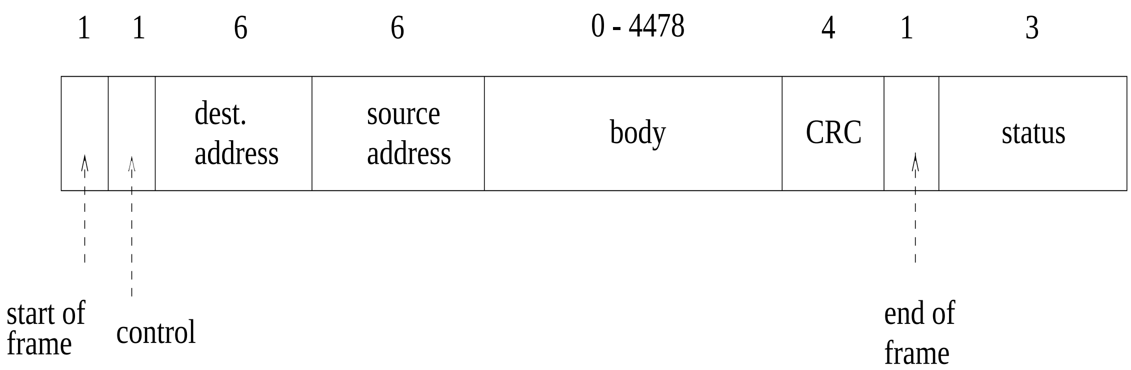

- frame size < 4500 B
- 4B/5B encoding
- synchronous/asynchronous data
- 2 km inter-station distance
- 200 km diameter (multimode fiber); 100 km circumference

Performance issues: fairness and efficiency

- TRT (token rotation time)
- THT (token holding time)

$$\text{TRT} = \text{no. of nodes} \times \text{THT} + \text{link latency}$$

To increase efficiency: increase THT

$\longrightarrow$ let station send as much as it needs

$\longrightarrow$ same as frame size $\uparrow$

$\longrightarrow$ THT $\uparrow \implies \rho \uparrow$

To increase fairness: limit THT

$\longrightarrow$ limit station's one-time sending of data

To facilitate fairness: introduce TTRT (target token rotation time).

THT determining factor (assume TTRT is given):

- prioritized frames: synchronous/asynchronous
- Synchronous frames always get sent.
- If TRT $>$ TTRT, then late; don't send asynchronous data.
- If TRT $\leq$ TTRT, then early; send asynchronous data for max $\{$ TTRT $-$ TRT, single frame time $\}$ duration.

How to set TTRT?

$\longrightarrow$ token claim process

$\longrightarrow$ initiate when needed (e.g., start-up)

- Each station submits claim frame containing TTRT bid.
- Smaller TTRT bid overrides higher TTRT bids.
  - Compare claim frame bid against own desired TTRT.
  - If less, then reset own TTRT to lower value.
  - If larger, then put lower bid on claim frame and forward.
- Winner: same bid value when claim frame makes full circle.

  $\longrightarrow$ leader election

At the end of the day, consistent TTRT value among all stations.

$\longrightarrow$ consensus problem

Compare against Ethernet's CSMA/CD.

⟶ round-robin reservation

⟶ absence of MA and collision

⟶ determinism vs. indeterminism

⟶ imperfect QoS assurance

⟶ performance vis-à-vis CSMA/CD?

Cooperative vs. noncooperative protocols

⟶ robust if some users use selfish MAC

⟶ could be malicious