

Real-World ARQ Performance: TCP



Ex.: Purdue \rightarrow UCSD

- Purdue (NSL): web server
- UCSD: web client

```
traceroute to planetlab3.ucsd.edu (132.239.17.226), 30 hops max, 40 byte packets
 1  switch-lwsn2133-z1r11 (128.10.27.250)  0.483 ms  0.344 ms  0.362 ms
 2  lwsn-b143-c6506-01-tcom (128.10.127.251)  0.488 ms  0.489 ms  0.489 ms
 3  172.19.57.1 (172.19.57.1)  0.486 ms  0.488 ms  0.489 ms
 4  tel-210-m10i-01-campus.tcom.purdue.edu (192.5.40.54)  0.614 ms  0.617 ms  0.615 ms
 5  gigapop.tcom.purdue.edu (192.5.40.134)  1.743 ms  1.679 ms  1.727 ms
 6  * * *
 7  * * *
 8  * * *
 9  hpr-lax-hpr--nlr-packetnet.cenic.net (137.164.26.130)  56.919 ms  56.919 ms  57.658 ms
10  hpr-ucsd-10ge--lax-hpr.cenic.net (137.164.27.165)  60.326 ms  60.198 ms  60.196 ms
11  nodeb-720--ucsd-t320-gw-10ge.ucsd.edu (132.239.255.132)  60.326 ms  60.370 ms  75.130 ms
```

\longrightarrow RTT \approx 60.5 msec

\longrightarrow receiver window size: 32 KB




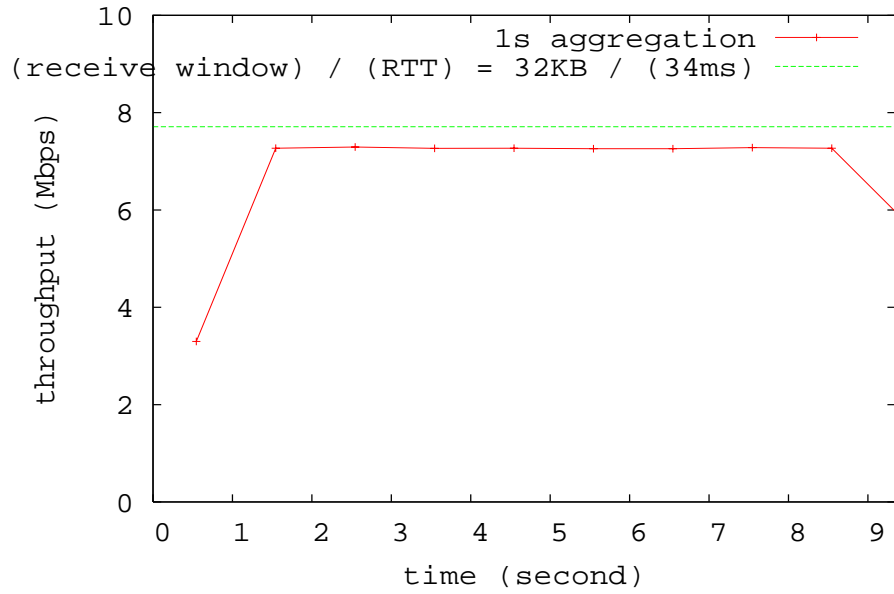
Ex.: Purdue \rightarrow Rutgers

- Purdue: web server
- Rutgers: web client

```
traceroute to planetlab1.rutgers.edu (165.230.49.114), 30 hops max, 40 byte packets
 1  switch-lwsn2133-z1r11 (128.10.27.250)  12.336 ms  0.339 ms  0.362 ms
 2  lwsn-b143-c6506-01-tcom (128.10.127.251)  0.489 ms  0.491 ms  0.488 ms
 3  172.19.57.1 (172.19.57.1)  0.490 ms  0.488 ms  0.489 ms
 4  tel-210-m10i-01-campus.tcom.purdue.edu (192.5.40.54)  0.614 ms  0.615 ms  0.614 ms
 5  switch-data.tcom.purdue.edu (192.5.40.166)  2.864 ms  2.865 ms  2.864 ms
 6  abilene-ul.indiana.gigapop.net (192.12.206.249)  2.988 ms  13.608 ms  3.113 ms
 7  chinng-iplsng.abilene.ucaid.edu (198.32.8.76)  6.740 ms  6.875 ms  6.859 ms
 8  ge-0-0-0.10.rtr.chic.net.internet2.edu (64.57.28.1)  7.113 ms  6.975 ms  6.986 ms
 9  so-3-0-0.0.rtr.wash.net.internet2.edu (64.57.28.13)  29.349 ms  24.086 ms  23.626 ms
10  ge-1-0-0.418.rtr.chic.net.internet2.edu (64.57.28.10)  44.786 ms  28.822 ms  28.839 ms
11  local.internet2.magpi.net (216.27.100.53)  30.723 ms  30.818 ms  30.744 ms
12  phl-02-09.backbone.magpi.net (216.27.100.229)  31.045 ms  36.644 ms  30.839 ms
13  remote.njedge.magpi.net (216.27.98.42)  33.221 ms  33.021 ms  33.087 ms
14  er01-hill-ext.runet.rutgers.net (198.151.130.233)  33.229 ms  33.207 ms  33.217 ms
```

\longrightarrow RTT \approx 34 msec

\longrightarrow receiver window size: 32 KB



Ex.: Purdue \rightarrow Korea University (Seoul)



- Purdue: web server
- KU: web client

```
1 switch-lwsn2133-z1r11 (128.10.27.250) 0.513 ms 10.061 ms 0.358 ms
2 lwsn-b143-c6506-01-tcom (128.10.127.251) 0.487 ms 0.476 ms 0.364 ms
3 172.19.57.1 (172.19.57.1) 0.489 ms 0.475 ms 0.490 ms
4 tel-210-m10i-01-campus.tcom.purdue.edu (192.5.40.54) 0.613 ms 0.600 ms 0.614 ms
5 switch-data.tcom.purdue.edu (192.5.40.166) 7.982 ms 7.969 ms 14.596 ms
6 abilene-ul.indiana.gigapop.net (192.12.206.249) 8.977 ms 7.721 ms 6.857 ms
7 kscyng-iplsng.abilene.ucaid.edu (198.32.8.81) 36.860 ms 25.873 ms 29.075 ms
8 dnvrng-kscyng.abilene.ucaid.edu (198.32.8.13) 24.218 ms 23.125 ms 36.317 ms
9 snvang-dnvrng.abilene.ucaid.edu (198.32.8.1) 47.815 ms 78.440 ms 54.048 ms
10 losang-snvang.abilene.ucaid.edu (198.32.8.94) 55.080 ms 55.131 ms 60.674 ms
11 transpac-1-lo-jmb-702.lsanca.pacificwave.net (207.231.240.136) 55.165 ms 55.212 ms 59.1
12 tokyo-losa-tp2.transpac2.net (192.203.116.146) 175.068 ms 170.832 ms 170.444 ms
13 tyo-gate1.jp.apan.net (203.181.248.249) 170.488 ms 170.893 ms 171.818 ms
14 sg-so-02-622m.bb-v4.noc.tein2.net (202.179.249.5) 277.150 ms 275.966 ms 276.136 ms
15 kr.pr-v4.noc.tein2.net (202.179.249.18) 278.422 ms 276.486 ms 280.132 ms
16 61.252.48.182 (61.252.48.182) 276.170 ms 279.606 ms 279.421 ms
17 202.30.43.45 (202.30.43.45) 271.663 ms 269.492 ms 268.761 ms
18 honeung13-seoul.kreonet.net (134.75.120.2) 269.781 ms 269.913 ms 273.516 ms
19 203.241.173.74 (203.241.173.74) 272.663 ms 278.774 ms 270.902 ms
```

\longrightarrow RTT \approx 292 msec

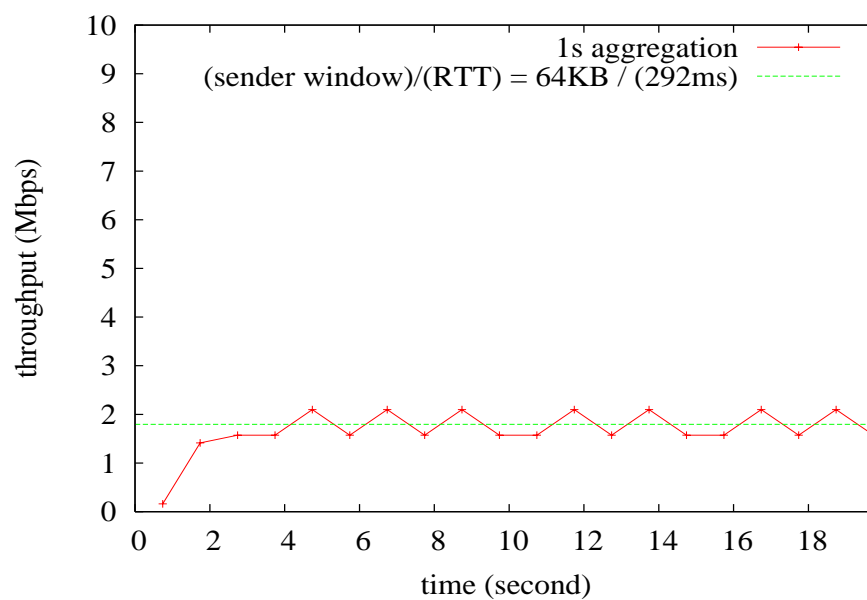
\longrightarrow long route to Korea (via Singapore)

\longrightarrow receiver window size: 32 KB



Increase receiver window size: 128 KB

→ 4-fold increase

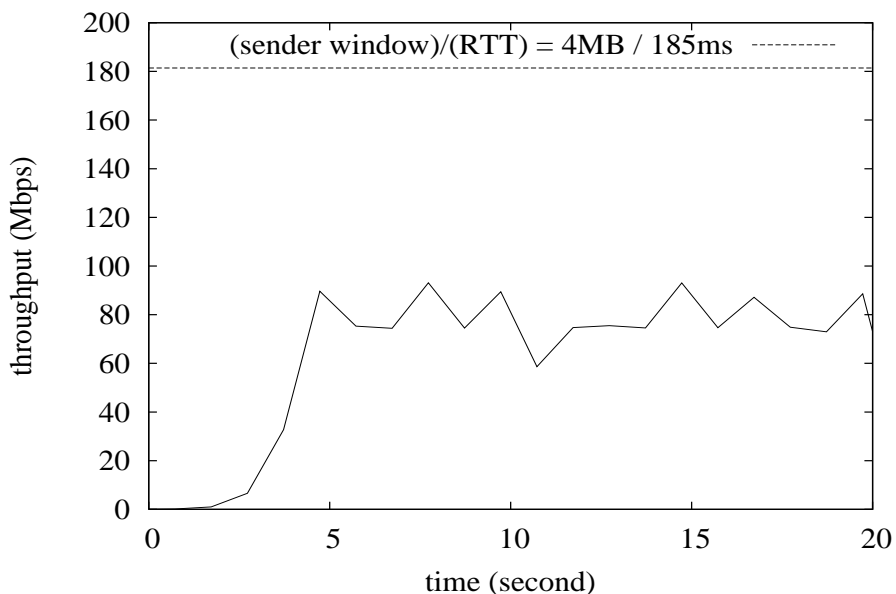


→ why only 2-fold throughput increase?

Increase receiver window size: 8 MB

→ also increase sender buffer size to 4 MB

→ $RTT \approx 185$ msec (short route to Korea)



→ around 90 Mbps

→ download time for 10 MB file?

→ can be confused with DoS (denial-of-service) attack

→ why less than 180 Mbps?

LINK LAYER: WIRED MEDIA

Multi-Access Communication

Bandwidth sharing: two approaches

- contention-free
 - e.g., TDM, FDM, TDMA, CDMA, token ring
 - used in telephony and broadband data networks
- contention-based
 - e.g., CSMA/CD, CSMA/CA
 - used in Ethernet, WLAN

... also called MAC (medium access control)

- broadband: FDM, TDMA, CDMA
- baseband: TDM, multiple access

Contention-free MAC:

- how to keep discussion orderly?
- moderator
- i.e., prior allocation of bandwidth



Contention-based MAC:

- single carrier frequency shared by many
- no moderator
- much less structured
- interference between conversations



Basic features:

- To send data, just send
- If two or more users send at the same time, data can become junk (i.e., bit flips)
 - called collision
 - collision need not always cause bit flips

Can get pretty “chaotic”

- used in real systems?
- yes, ALOHA (1970s)
- packet network connecting Univ. of Hawaii campuses
- ~30 years before boom of wireless networks!
- today’s technology: hasn’t changed much

Additional features (optional):

- NIC can detect if someone else is using the channel
 - called carrier sense (CS)
 - rule: if someone else talks, don't talk
 - impose gentlemanly (i.e., cooperative) behavior
- NIC can detect if collision has occurred
 - called collision detection (CD)
 - not always possible

In real systems:

- Ethernet: CS + CD
- WLAN: CS

Why not just use TDM?

Benefits of contention-based MAC:

- when not too many users, faster response time
 - minimal coordination overhead
 - e.g.: in TDM need to request and reserve slots
- decentralized
 - minimal configuration overhead
 - to join: just send
 - except for security concerns (e.g., Purdue's PAL)
 - a good idea?

Drawbacks of contention-based MAC:

- when many users, degraded throughput
 - collision wastes slots, i.e., bandwidth
- lack of QoS (quality of service) assurances
 - “you get what you get”; called best effort
 - problematic for real-time traffic, e.g., telephony

When to use what?

→ contention-free vs. contention-based

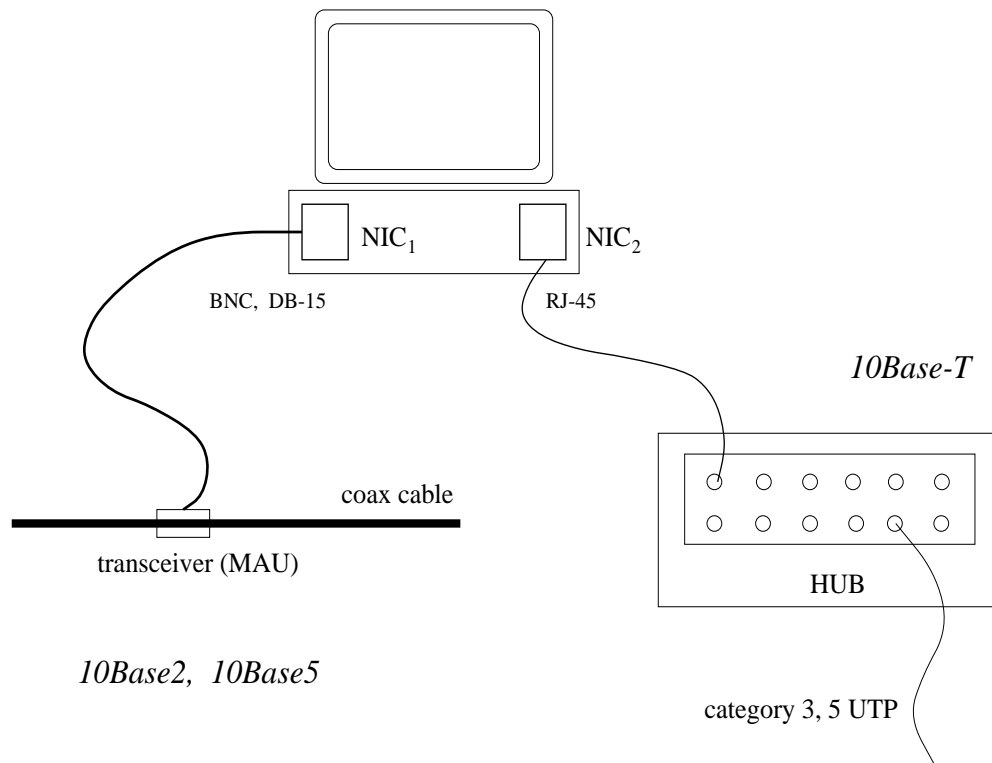
Ethernet

→ copper, fiber

Types:

- 10Base2 (ThinNet): coax, segment length 200 m, 30 nodes/segment
- 10Base5 (ThickNet): coax, segment length 500 m, 100 nodes/segment
- 10Base-T: twisted pair, segment length 100 m, 1024 nodes/segment
- 100Base-T (Fast Ethernet): category 5 UTP, fiber (also 100VG-AnyLAN)
- Gigabit & 10 Gbps Ethernet: fiber, category 5 UTP
- 100 Gbps Ethernet: soon

Connectivity example:




- single-homed vs. multi-homed
- unique 48-bit Ethernet address per NIC
- physical network: bus vs. hub vs. switch
→ very old vs. old vs. not-so-old

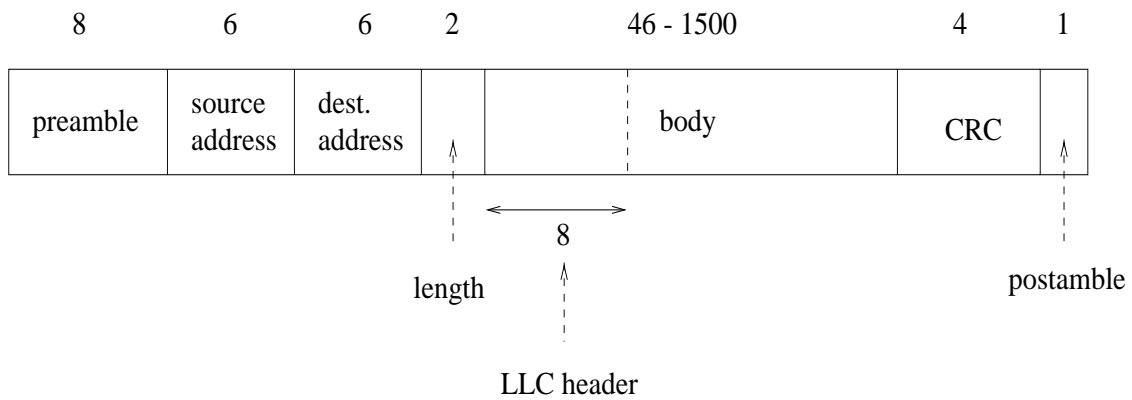
High-speed Ethernets have shorter network diameter

- about 2500 m for 10 Mbps Ethernet
 - about 200 m for 100 Mbps Ethernet
 - even shorter for 1 Gbps Ethernet
-
- distance limitations: due to Ethernet protocol
 - creates complications for long-haul
 - i.e., tens, hundreds, or thousands of miles
 - 1 and 10 Gbps: popular tier-1 backbone speeds

DIX Ethernet frame:



IEEE 802.3 Ethernet frame:



→ IEEE 802.2 LLC (Logical Link Control)

→ two Ethernet types co-exist

Encoding: Manchester

→ note: Ethernet is baseband

Addressing:

- 48 bit unique address
- point-to-point
- broadcast (all 1's)

Receiver: Ethernet NIC accepts frames with “relevant” address.

- accepts only own frame address
 - default
- accepts all frames: called promiscuous mode
 - can set with root privilege
 - useful for traffic monitoring/sniffing

Ethernet MAC protocol: CSMA/CD

- MA (Multiple Access): multiple nodes are allowed simultaneous access
 - just send
- CS (Carrier Sense): can detect if some other node is using the link
 - rule: if busy, wait until channel is not busy
- CD (Collision Detection): can detect if collision due has occurred
 - rule: if collision, retry later
 - key question: when is “later”?


Wired (Ethernet) vs. wireless (WLAN):

- CD is key difference
- wireless: difficult to detect collision

Signal propagation and collision:

Bi-directional propagation


→ terminator absorbs signal: prevent bounce back



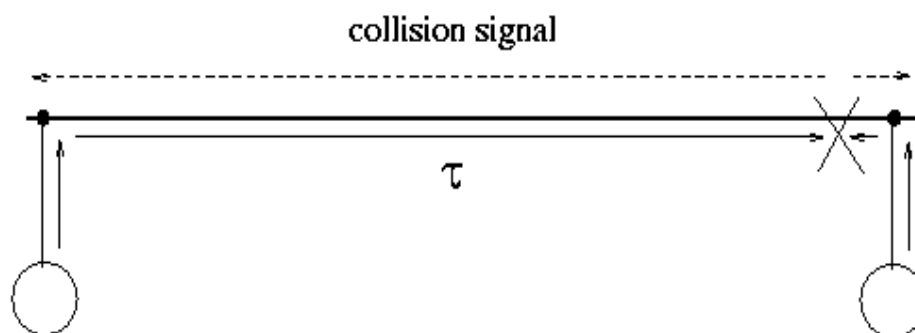
Best-case collision: 2 stations

→ meet in the middle: why?

→ what is the worst-case?



Worst-case collision scenario:



→ τ : one-way propagation delay

- sender needs to wait 2τ sec before detecting collision
→ time for echo to bounce back
- for 2500 m length, $51.2 \mu\text{s}$ round-trip time (2τ)
- enforce $51.2 \mu\text{s}$ slot time
- at 10 Mbps, 512 bits; i.e., minimum frame size
→ assures collision detection
→ wireless: bounceback problem

Transmit at least 512 bits

$$\longrightarrow 6 + 6 + 2 + 46 + 4 = 64 \text{ B} = 512 \text{ bits}$$

When to retry upon collision: use exponential backoff

1. Wait for random $0 \leq X \leq 51.2 \mu\text{s}$ before 1st retry
2. Two consecutive collisions: wait for random $0 \leq X \leq 102.4 \mu\text{s}$ before 2nd retry
3. Three consecutive collisions: wait for random $0 \leq X \leq 204.8 \mu\text{s}$ before 3rd retry
2. i consecutive collisions: wait for $0 \leq X \leq 2^{i-1} 51.2 \mu\text{s}$ before next attempt
3. Give up if $i > 16$

\longrightarrow a form of stop-and-wait

\longrightarrow what's the ACK?

\longrightarrow guaranteed reliability?

\longrightarrow why exponential backoff?