

LINK LAYER: BASIC TECHNIQUES

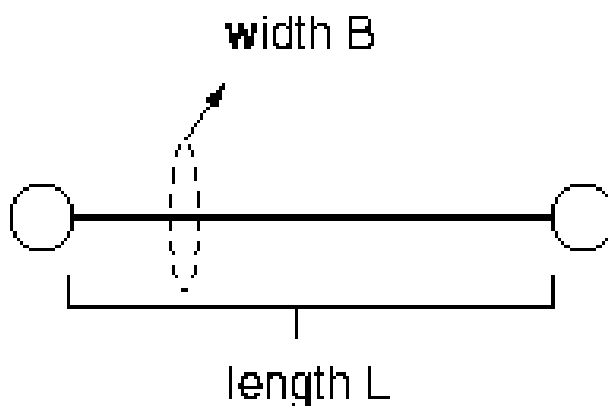
Data Transmission

Link speed unit: bps

→ abstraction

→ ignore carrier frequency, coding etc.

Simplest case: point-to-point link



→ wired or wireless

Interested in *completion time*:

→ time elapsed between sending/receiving first bit

→ i.e., how long will it take?

- Single bit:

→ $\approx L/SOL$ (lower bound)

→ latency (or propagation delay)

→ optical fiber, wireless: exact

- Multiple, say S , bits:

→ $\approx L/SOL + S/B$

→ latency + transmission time

Latency vs. transmission time: which dominates?

→ a lot to send, a little to send, ...

→ satellite, Zigbee, WLAN, broadband WAN

Reliable Transmission

Main method: ARQ (Automatic Repeat reQuest)

→ use retransmission

→ used in both wired/wireless

- function duplication

→ link layer, transport layer, etc.

- alternative: FEC (forward error correction)


→ transmit redundant information

→ not assured

→ pros and cons?


ARQ: three components

- timer
- acknowledgment (ACK)
- retransmit



Special case: stop-and-wait

Handle one packet (i.e., frame) at a time.



Issue of RTT (Round-Trip Time) & timer management:

- what is proper value of timer?
 - RTT estimation
- easier for single link
 - RTT is more well-behaved
- more difficult for multi-hop path in internet network
 - latency + queueing effect

A “good” thing about stop-and-wait:

→ simple throughput formula

Stop-and-wait throughput (bps):

- RTT
- frame size (bits)

$$\longrightarrow \text{throughput} = \text{frame size} / \text{RTT}$$

Another important problem: not keeping the pipe full.

→ delay-bandwidth product

→ volume of data travelling on the link

High throughput: want to keep the pipe full

Ex.: Link BW 1.5 Mbps, 45 msec RTT

- if frame size 1 kB, then throughput:
 - $1024 \times 8 / 0.045 = 182 \text{ kbps}$
 - utilization: only $182 \text{ kbps} / 1500 \text{ kbps} = 0.121$
- note: delay-bandwidth product
 - $1.5 \text{ Mbps} \times 45 \text{ msec} = 67.5 \text{ kb} \approx 8 \text{ kB}$

What happens to utilization if RTT increases to 90 msec?

What happens if bandwidth increases to 3 Mbps?


→ how to reduce bandwidth wastage?

Sliding Window Protocol

→ send block (i.e., window) of data

Issues:

- Shield application process from reliability management chore
 - exported semantics: continuous data stream
 - simple app abstraction: e.g., **read** system call
- Both sender and receiver have limited buffer capacity
 - task: plug holes & flush buffer



Simple solution when receiver has infinite buffer capacity:


- sender keeps sending at maximum speed
- receiver informs sender of holes
 - “I’m missing this and that”
 - called negative ACK
- sender retransmits missing frames

Drawbacks?

What about positive ACK?

→ pros and cons

Sliding window operation with positive ACK:



- *SWS*: Sender Window Size (sender buffer size)
- *RWS*: Receiver Window Size (receiver buffer size)
- *LAR*: Last ACK Received
- *LFS*: Last Frame Sent
- *NFE*: Next Frame Expected
- *LFA*: Last Frame Acceptable

Assign sequence numbers to frames.

→ IDs

Maintain invariants:

- $LFA - NFE + 1 \leq RWS$
- $LFS - LAR + 1 \leq SWS$

Sender:

- Receive ACK with sequence number X
- Forwind LAR to X
- Flush buffer up to (but not including) LAR
- Send up to $SWS - (LFS - LAR + 1)$ frames
- Update LFS

Receiver:

- Receive packet with sequence number Y
- Forward to (new) first hole & update NFE
→ NFE need not be $Y + 1$
- Send cumulative ACK (i.e., NFE)
- Flush buffer up to (but not including) NFE to application
- Update $LFA \leftarrow NFE + RWS - 1$

Sequence number wrap-around problem:

$$\text{SWS} < (\text{MaxSeqNum} + 1)/2$$

- why?
- consider special case: stop-and-wait
- is sequence number needed?

Real-World ARQ Performance: TCP


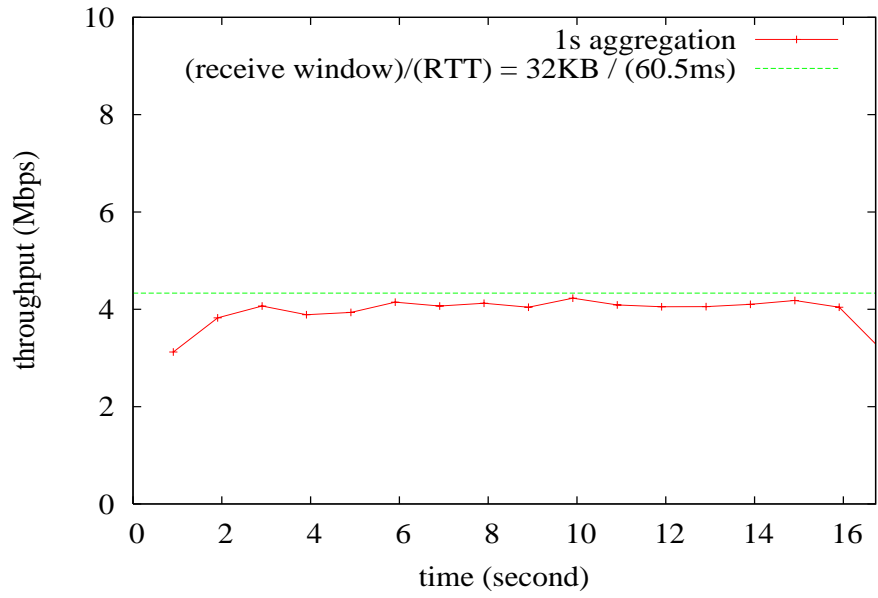
Ex.: Purdue \rightarrow UCSD

- 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
```

\rightarrow RTT \approx 60.5 msec

\rightarrow 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)


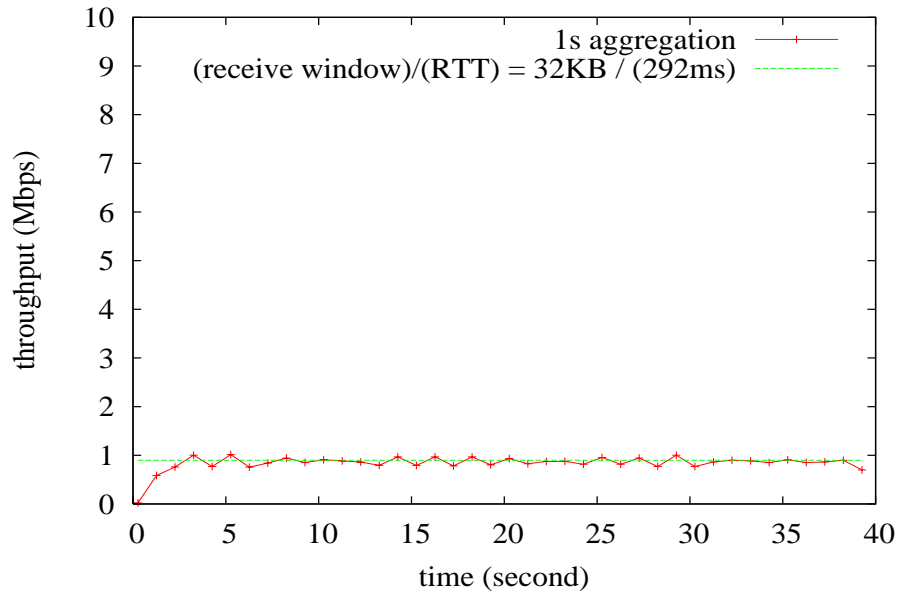
- Purdue: web server
- KU: web client

```
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?