

CS 725/825 & IT 725

Lecture 14

Transport Layer


October 16, 2024

Principles of Reliable Transport

- ▶ **Goal:** deliver despite unreliability of the network layer or detect that delivery is not possible
- ▶ **Automatic Repeat reQuest (ARQ):**
 - acknowledgment
 - timeout
 - retransmission
 - give up after k retransmissions
 - sequence numbers on data packets
 - cumulative acknowledgment numbers


Cumulative ACK numbers

- ▶ Method 1 (obvious but not used):
 - ACK carries the sequence number of the packet it acknowledges
- ▶ Method 2 (**Cumulative ACK**, used by TCP)
 - ACK carries the lowest sequence number of the packets that were not yet received (sequence number of the next expected packet)




CUMULATIVE ACKS

NOTATION



BENEFITS (EX.)



CONFIRMS
 7, 8, 9 →
 7, 8, 9 rec'd
 no need
 to retr.

Filling the pipe...

- ▶ **Stop and Wait** protocol
 - wait for acknowledgment before sending next packet
- ▶ **Sliding Window** protocols
 - send up to w (window size) packets/bytes before waiting for acknowledgment
 - when a packet is lost:
 - retransmit the packet (*Selective-Reject ARQ*)
 - retransmit all un-acknowledged packets (*Go-Back-N ARQ*)
- ▶ Measure: **utilization** (a.k.a. normalized throughput)
 - the ratio between *goodput* and *maximum theoretical capacity*

Flow & Congestion Control

▶ Receiver Congestion

- receiver is unable to keep up with incoming data
- solved by explicit feedback from receiver to sender

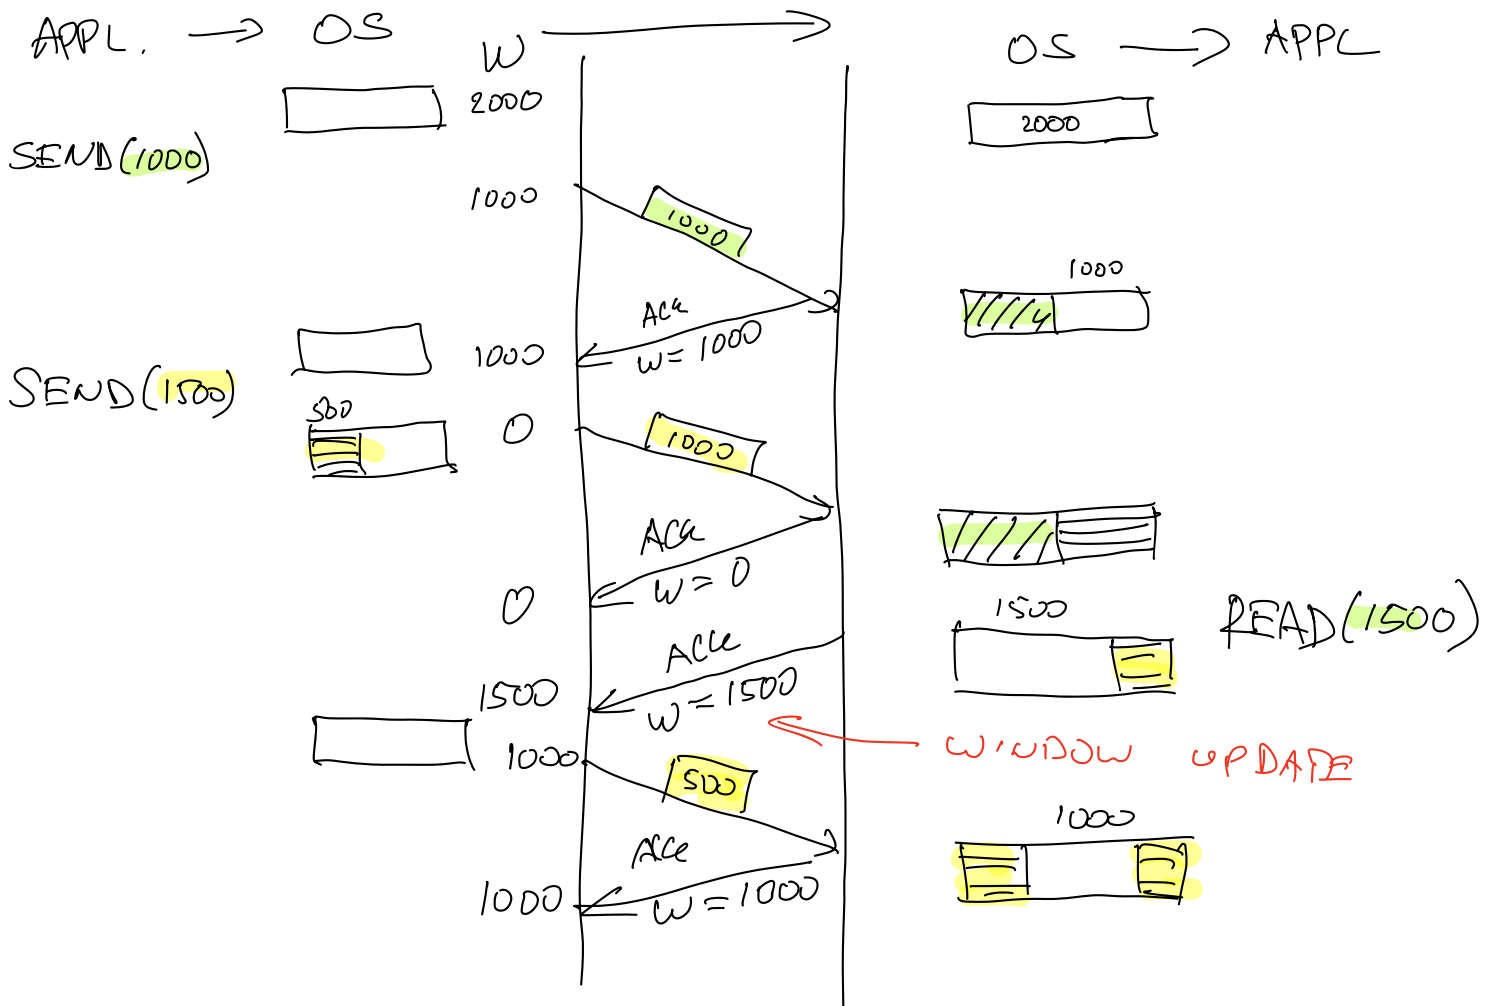
▶ Network Congestion

- nodes or links of the network are overloaded
- **explicit** congestion notification (few technologies)
- **implicit** congestion notification (Internet)

Congestion control

- ▶ **Goal:** Make the most effective use of the network capacity
 - avoid congestion
 - maximize utilization
 - maintain fairness (or deliver promised service level)
- ▶ **Method:** Controlling the rate with which traffic is injected into the network by the transmitter

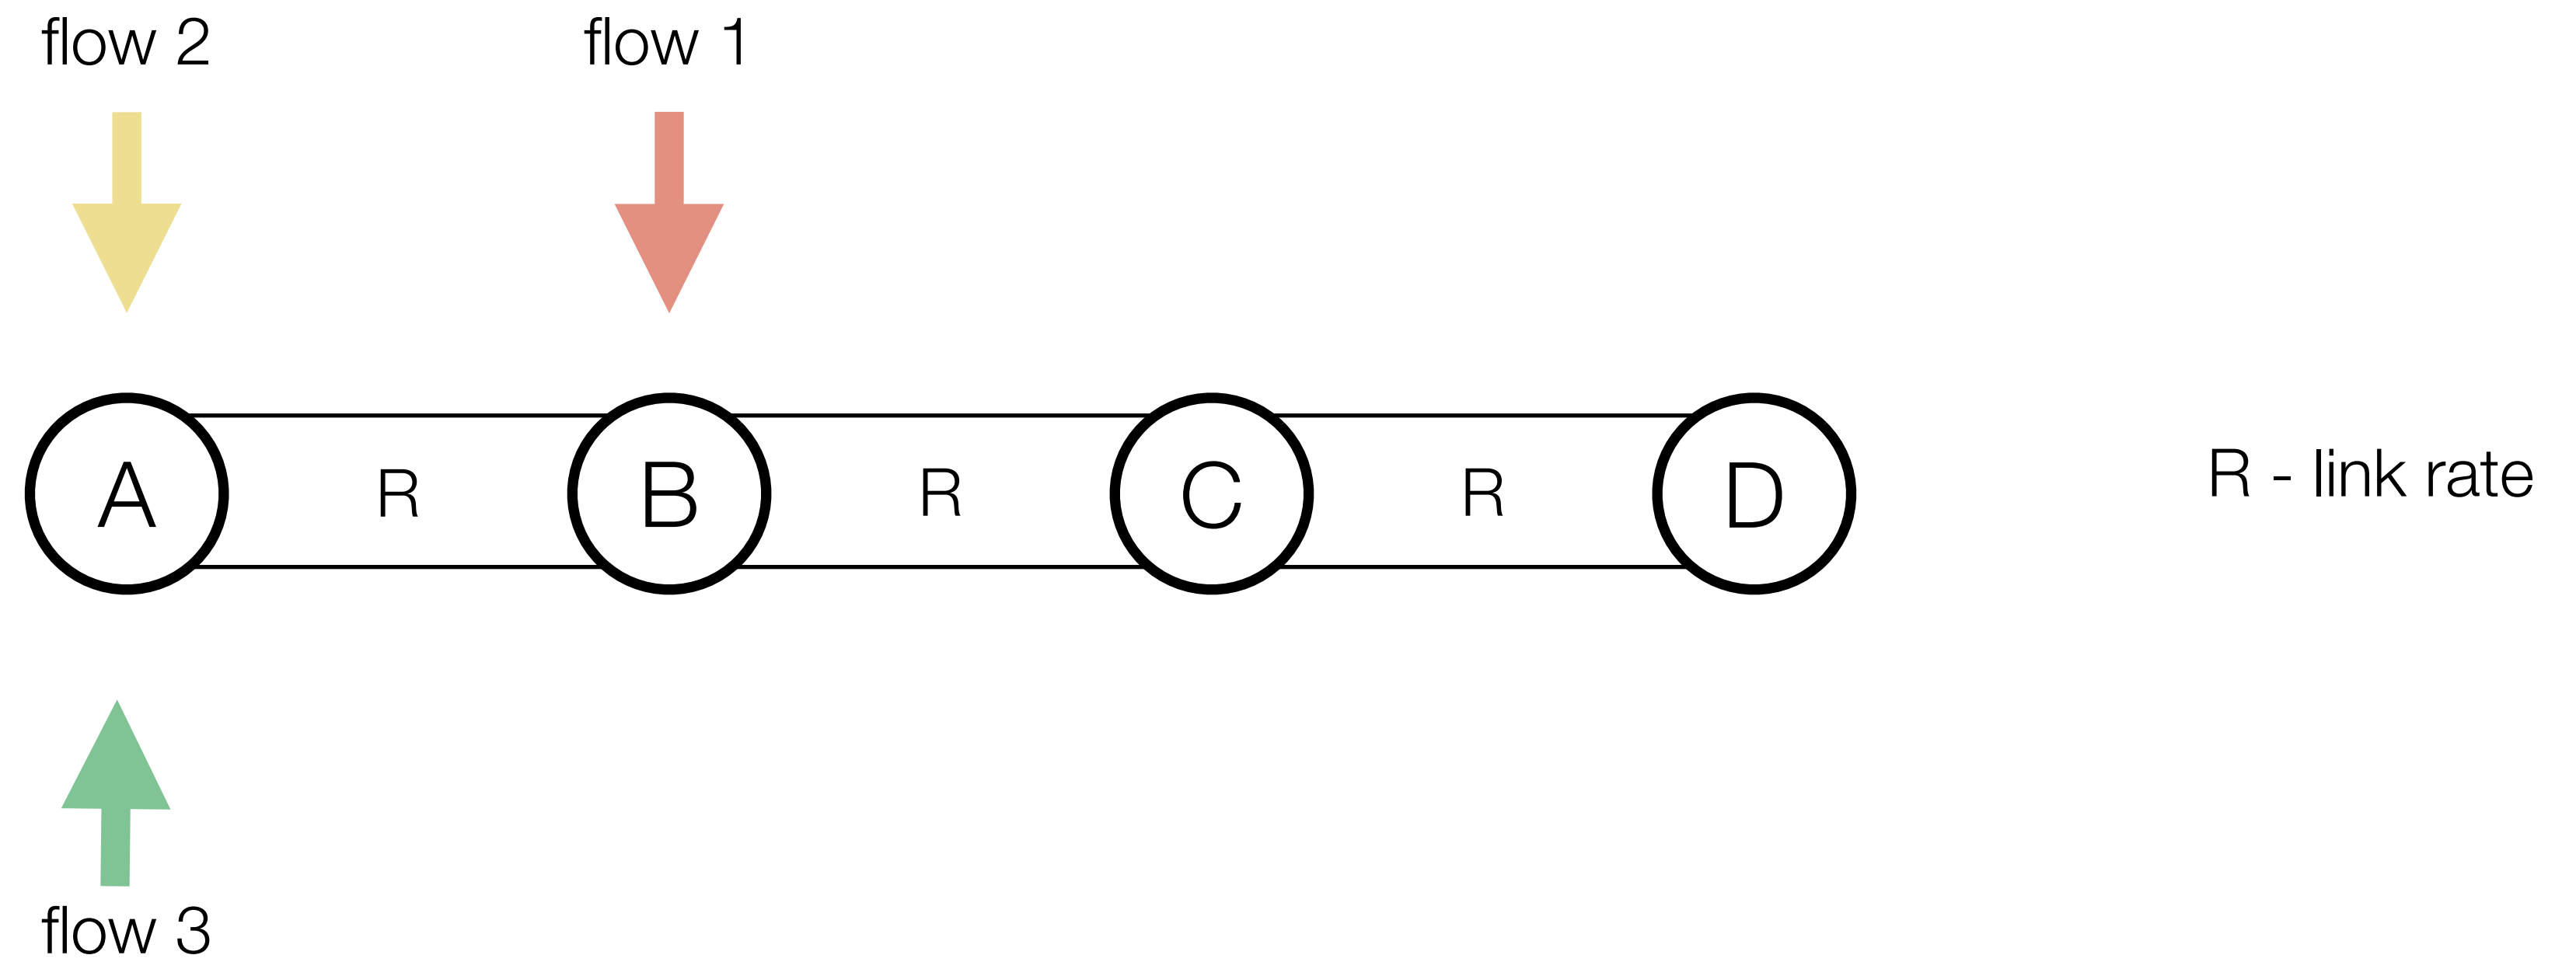
FLOW CONTROL



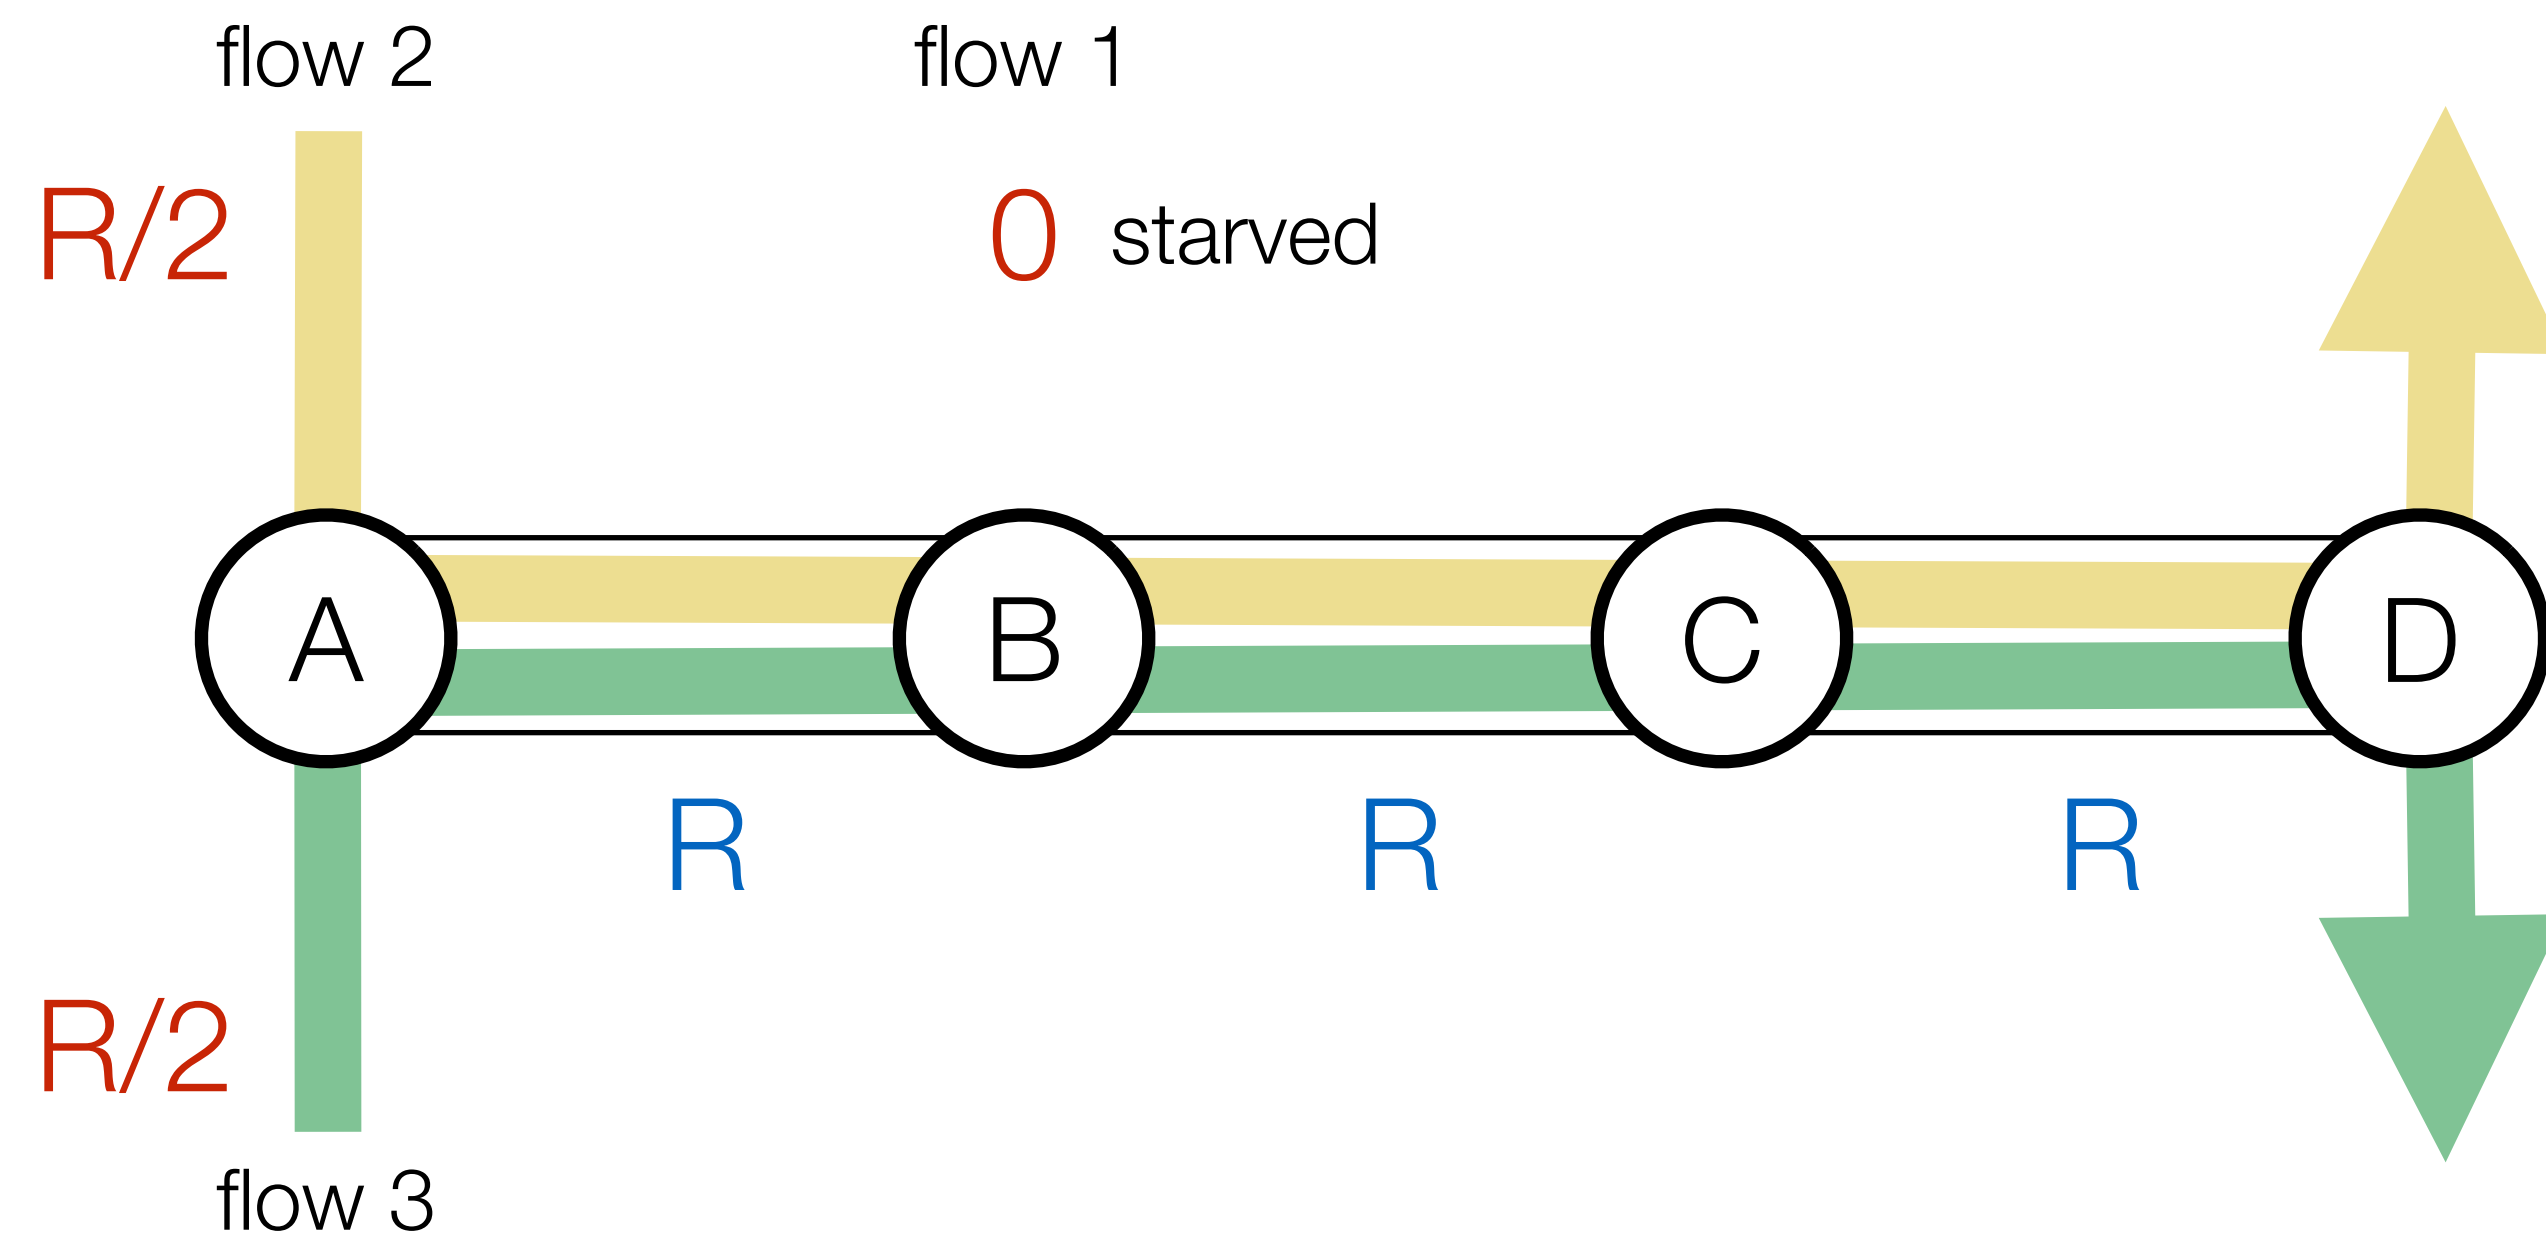
Congestion control

- ▶ Reasons why congestion control mechanisms are critical for the stable operation of the Internet [RFC 8085]:
- ▶ **Prevention of congestion collapse**
 - i.e., a state where an increase in network load results in a decrease in useful work done by the network
- ▶ **Establishment of a degree of fairness**
 - i.e., allowing multiple flows to share the capacity of a path reasonably equitably.

Utilization vs fairness



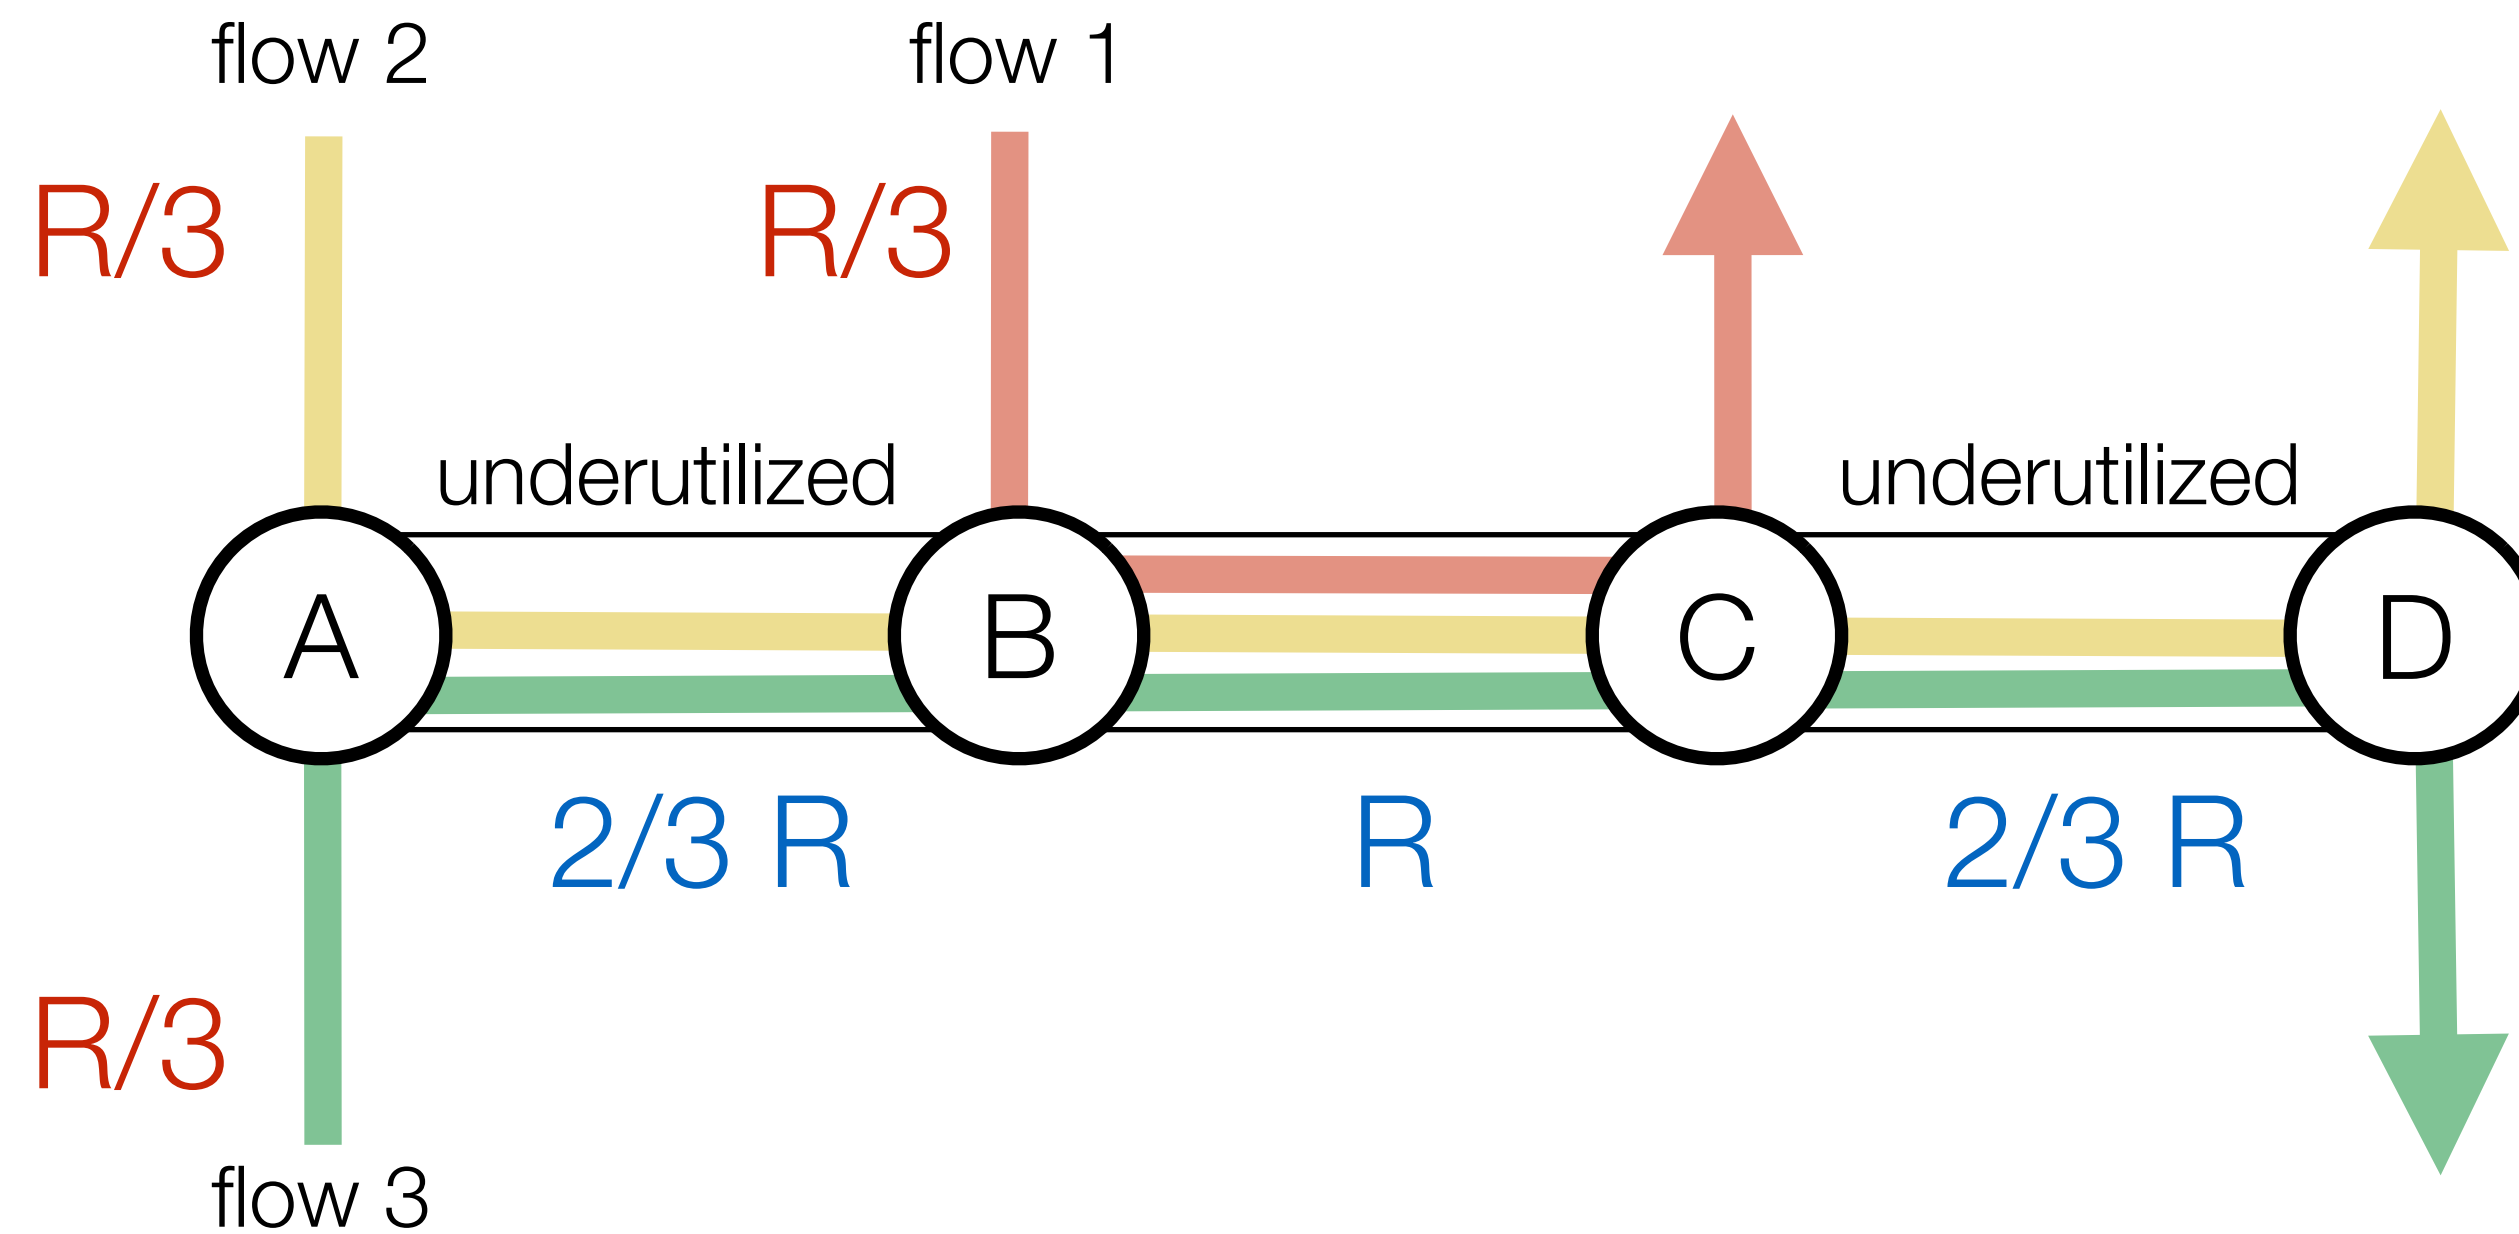
Utilization vs fairness



	Max utilization	Max fairness
flow 1	0	$R/3$
flow 2	$R/2$	$R/3$
flow 3	$R/2$	$R/3$

R - link rate

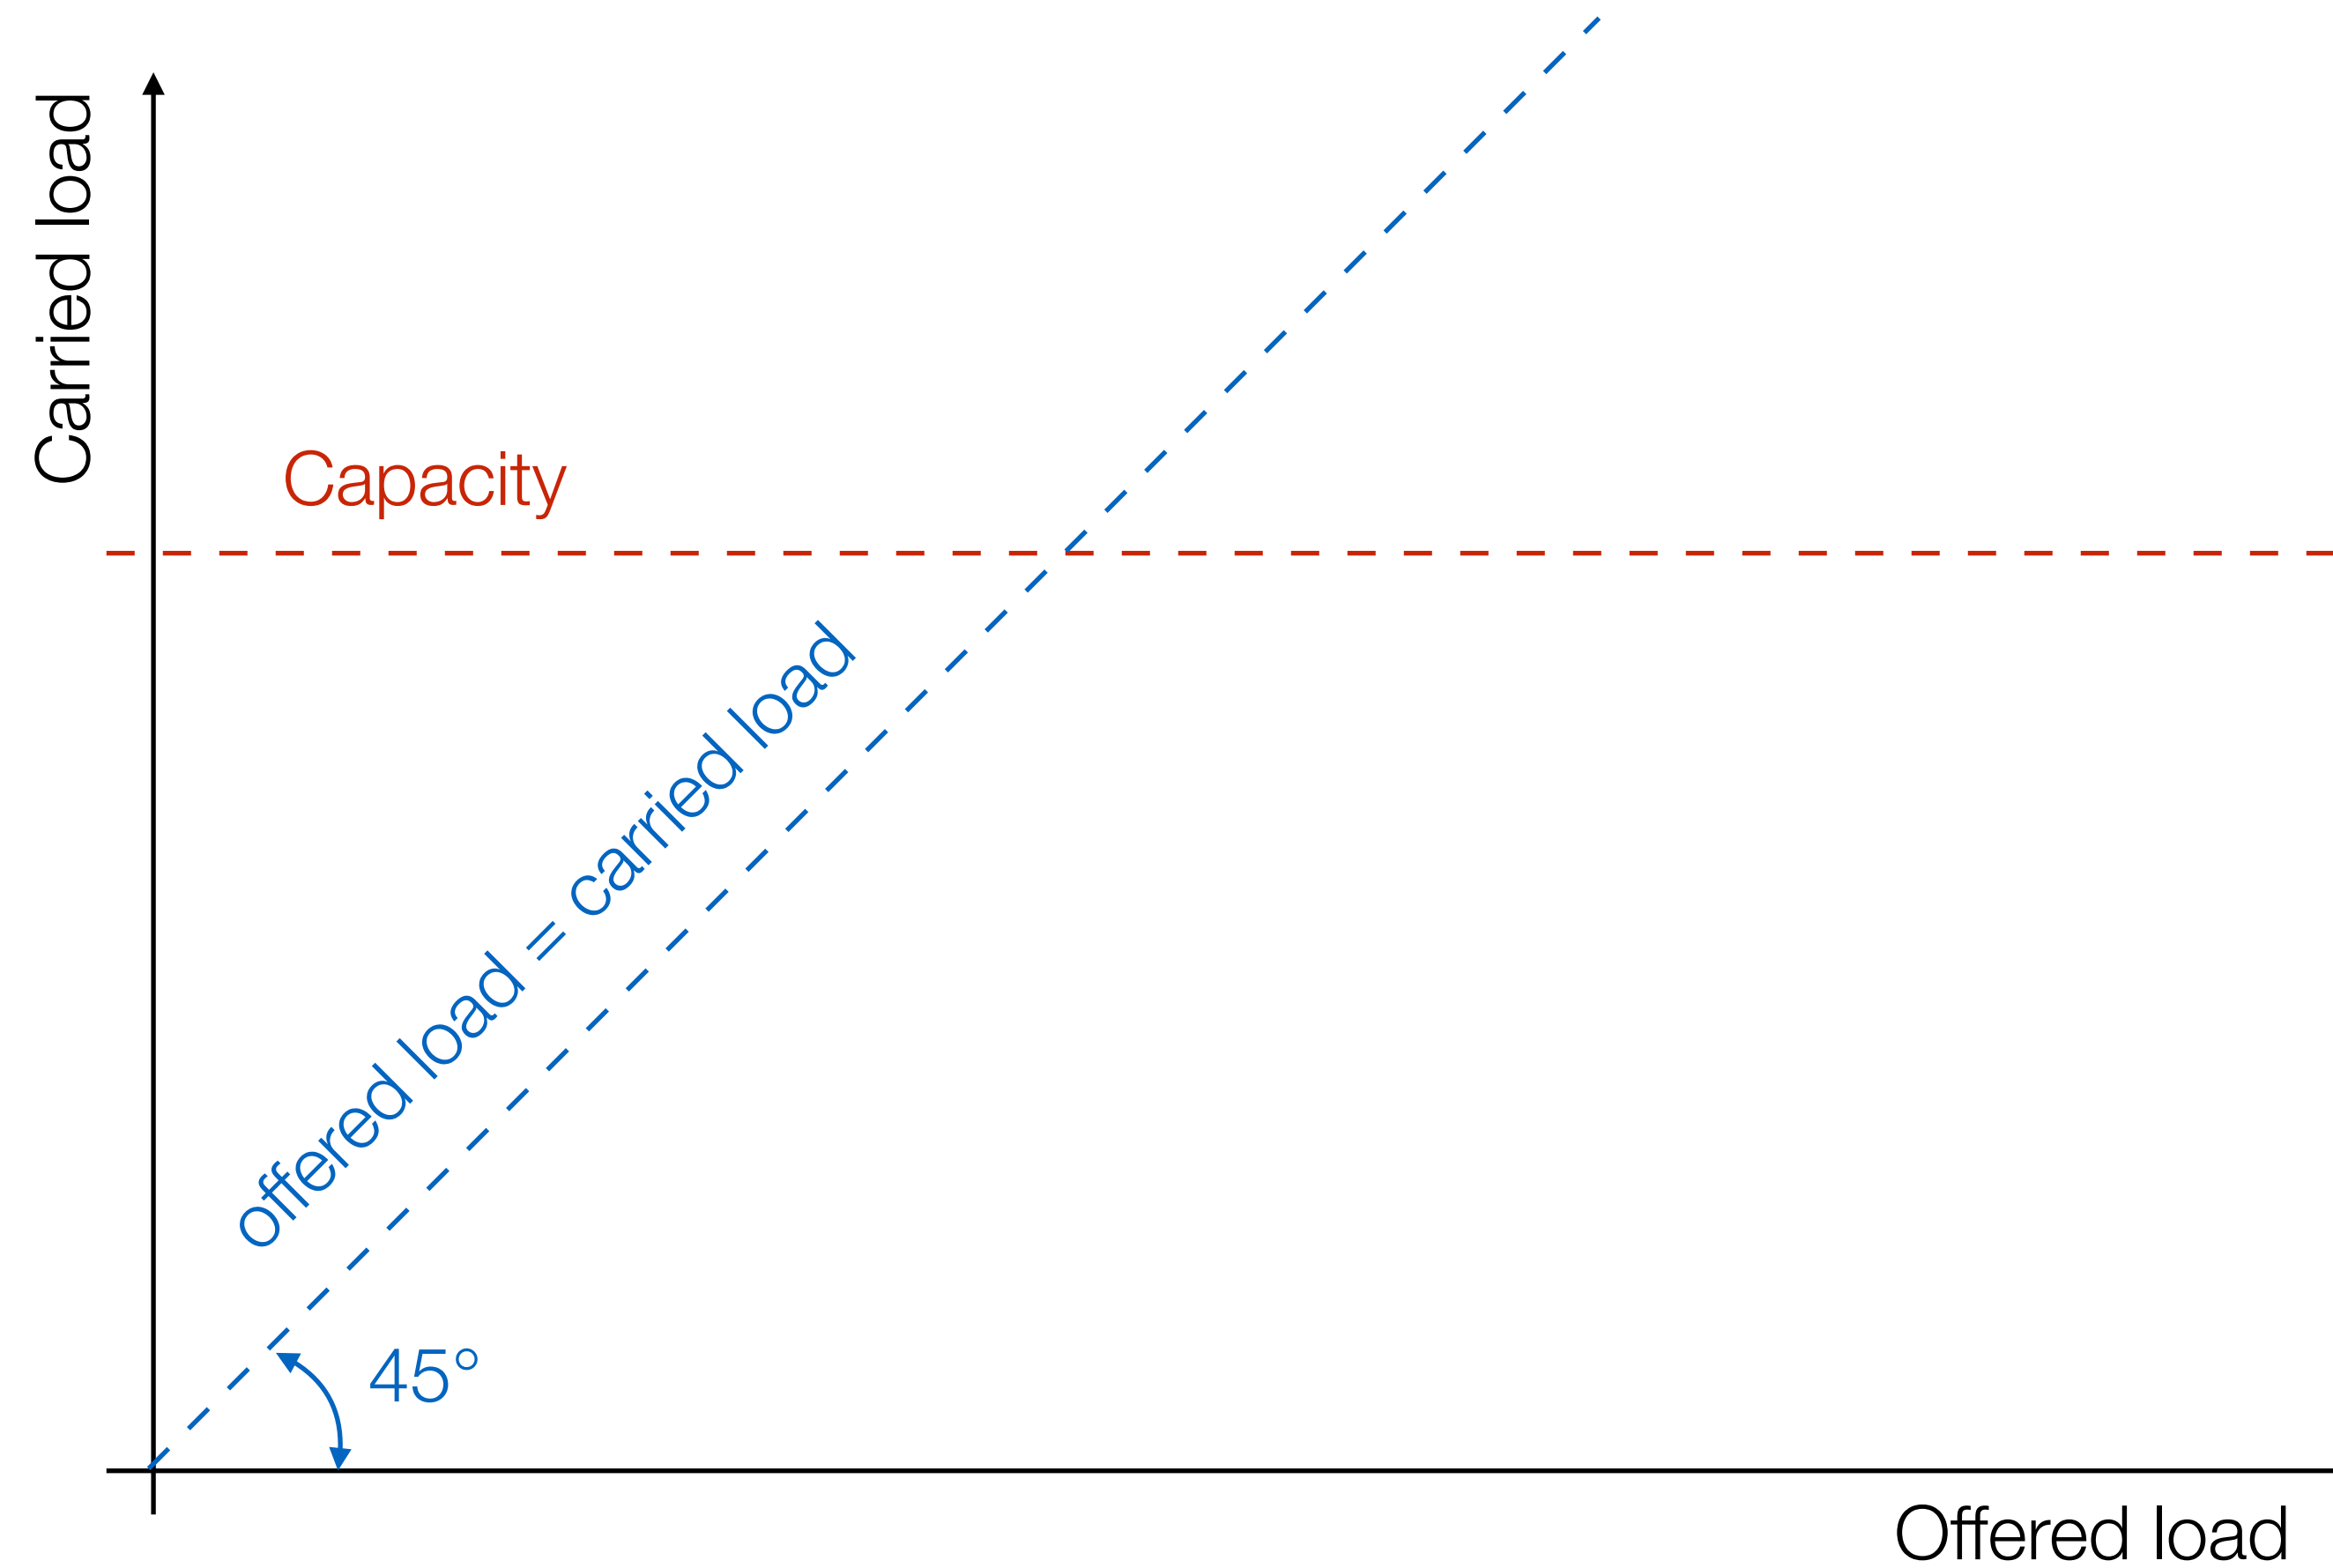
Utilization vs fairness



	Max utilization	Max fairness
flow 1	0	$R/3$
flow 2	$R/2$	$R/3$
flow 3	$R/2$	$R/3$


R - link rate

Impact of Congestion

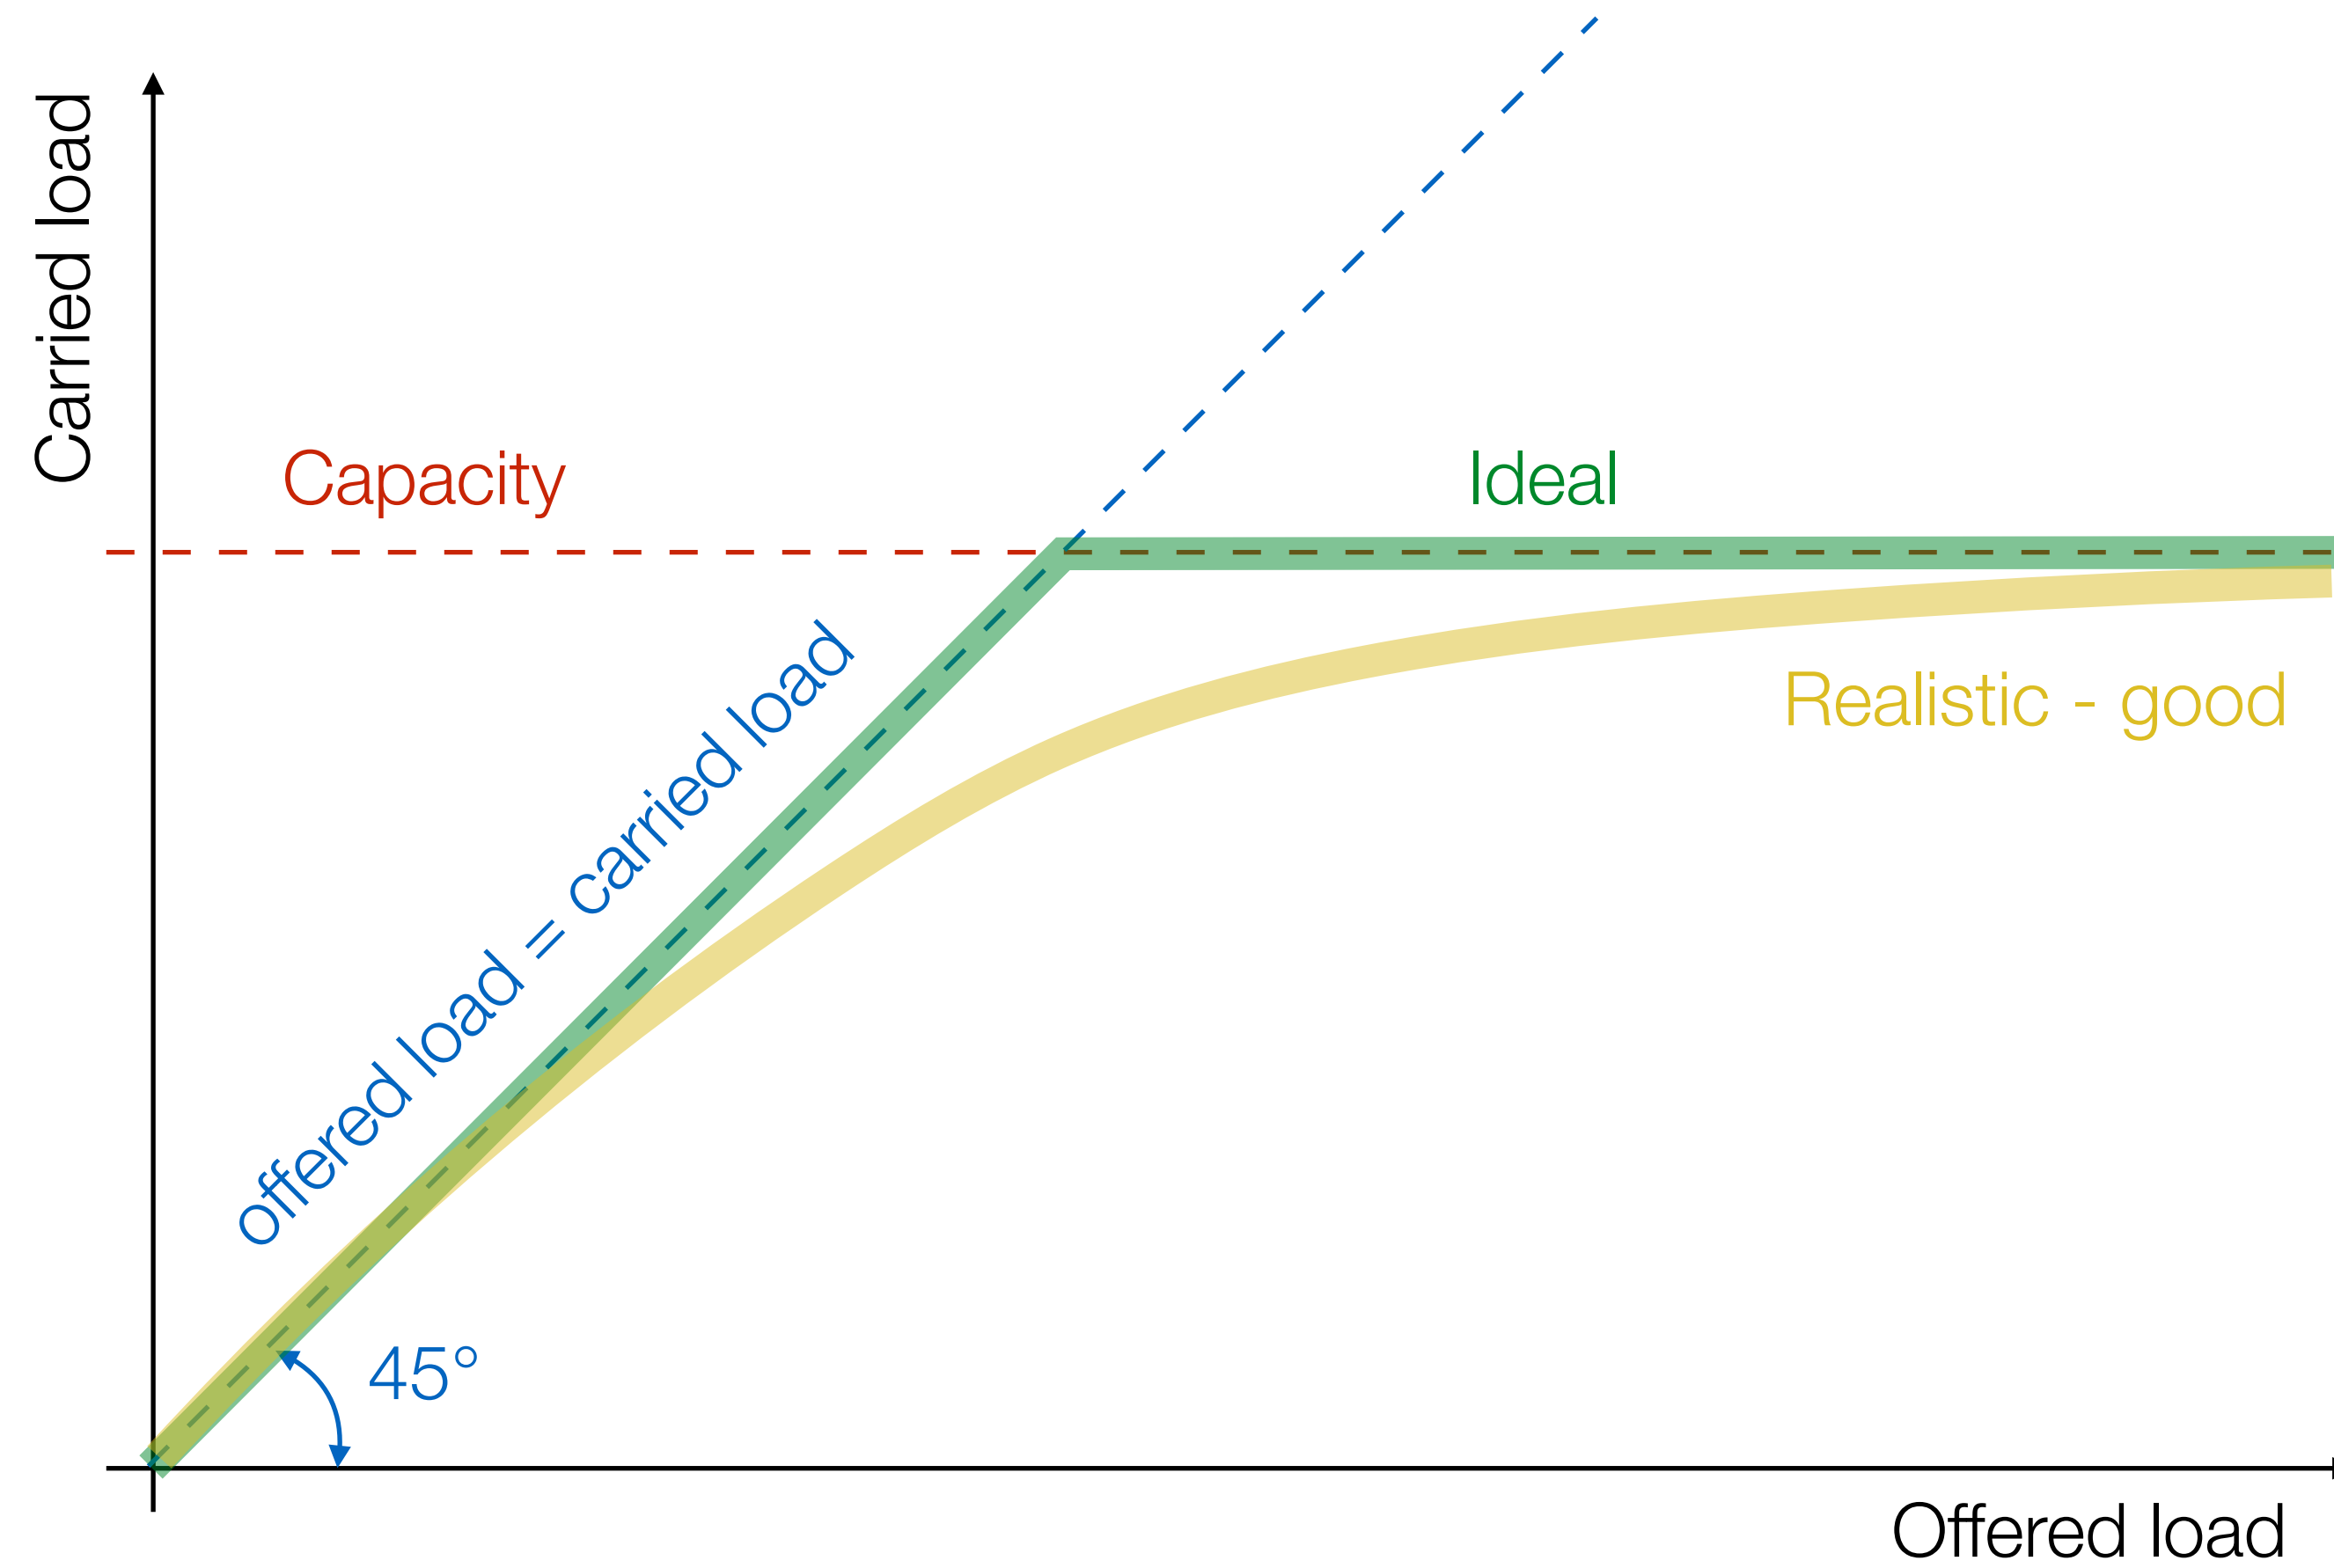


Offered vs carried load graph


Impact of Congestion



Impact of Congestion



Impact of Congestion



TCP

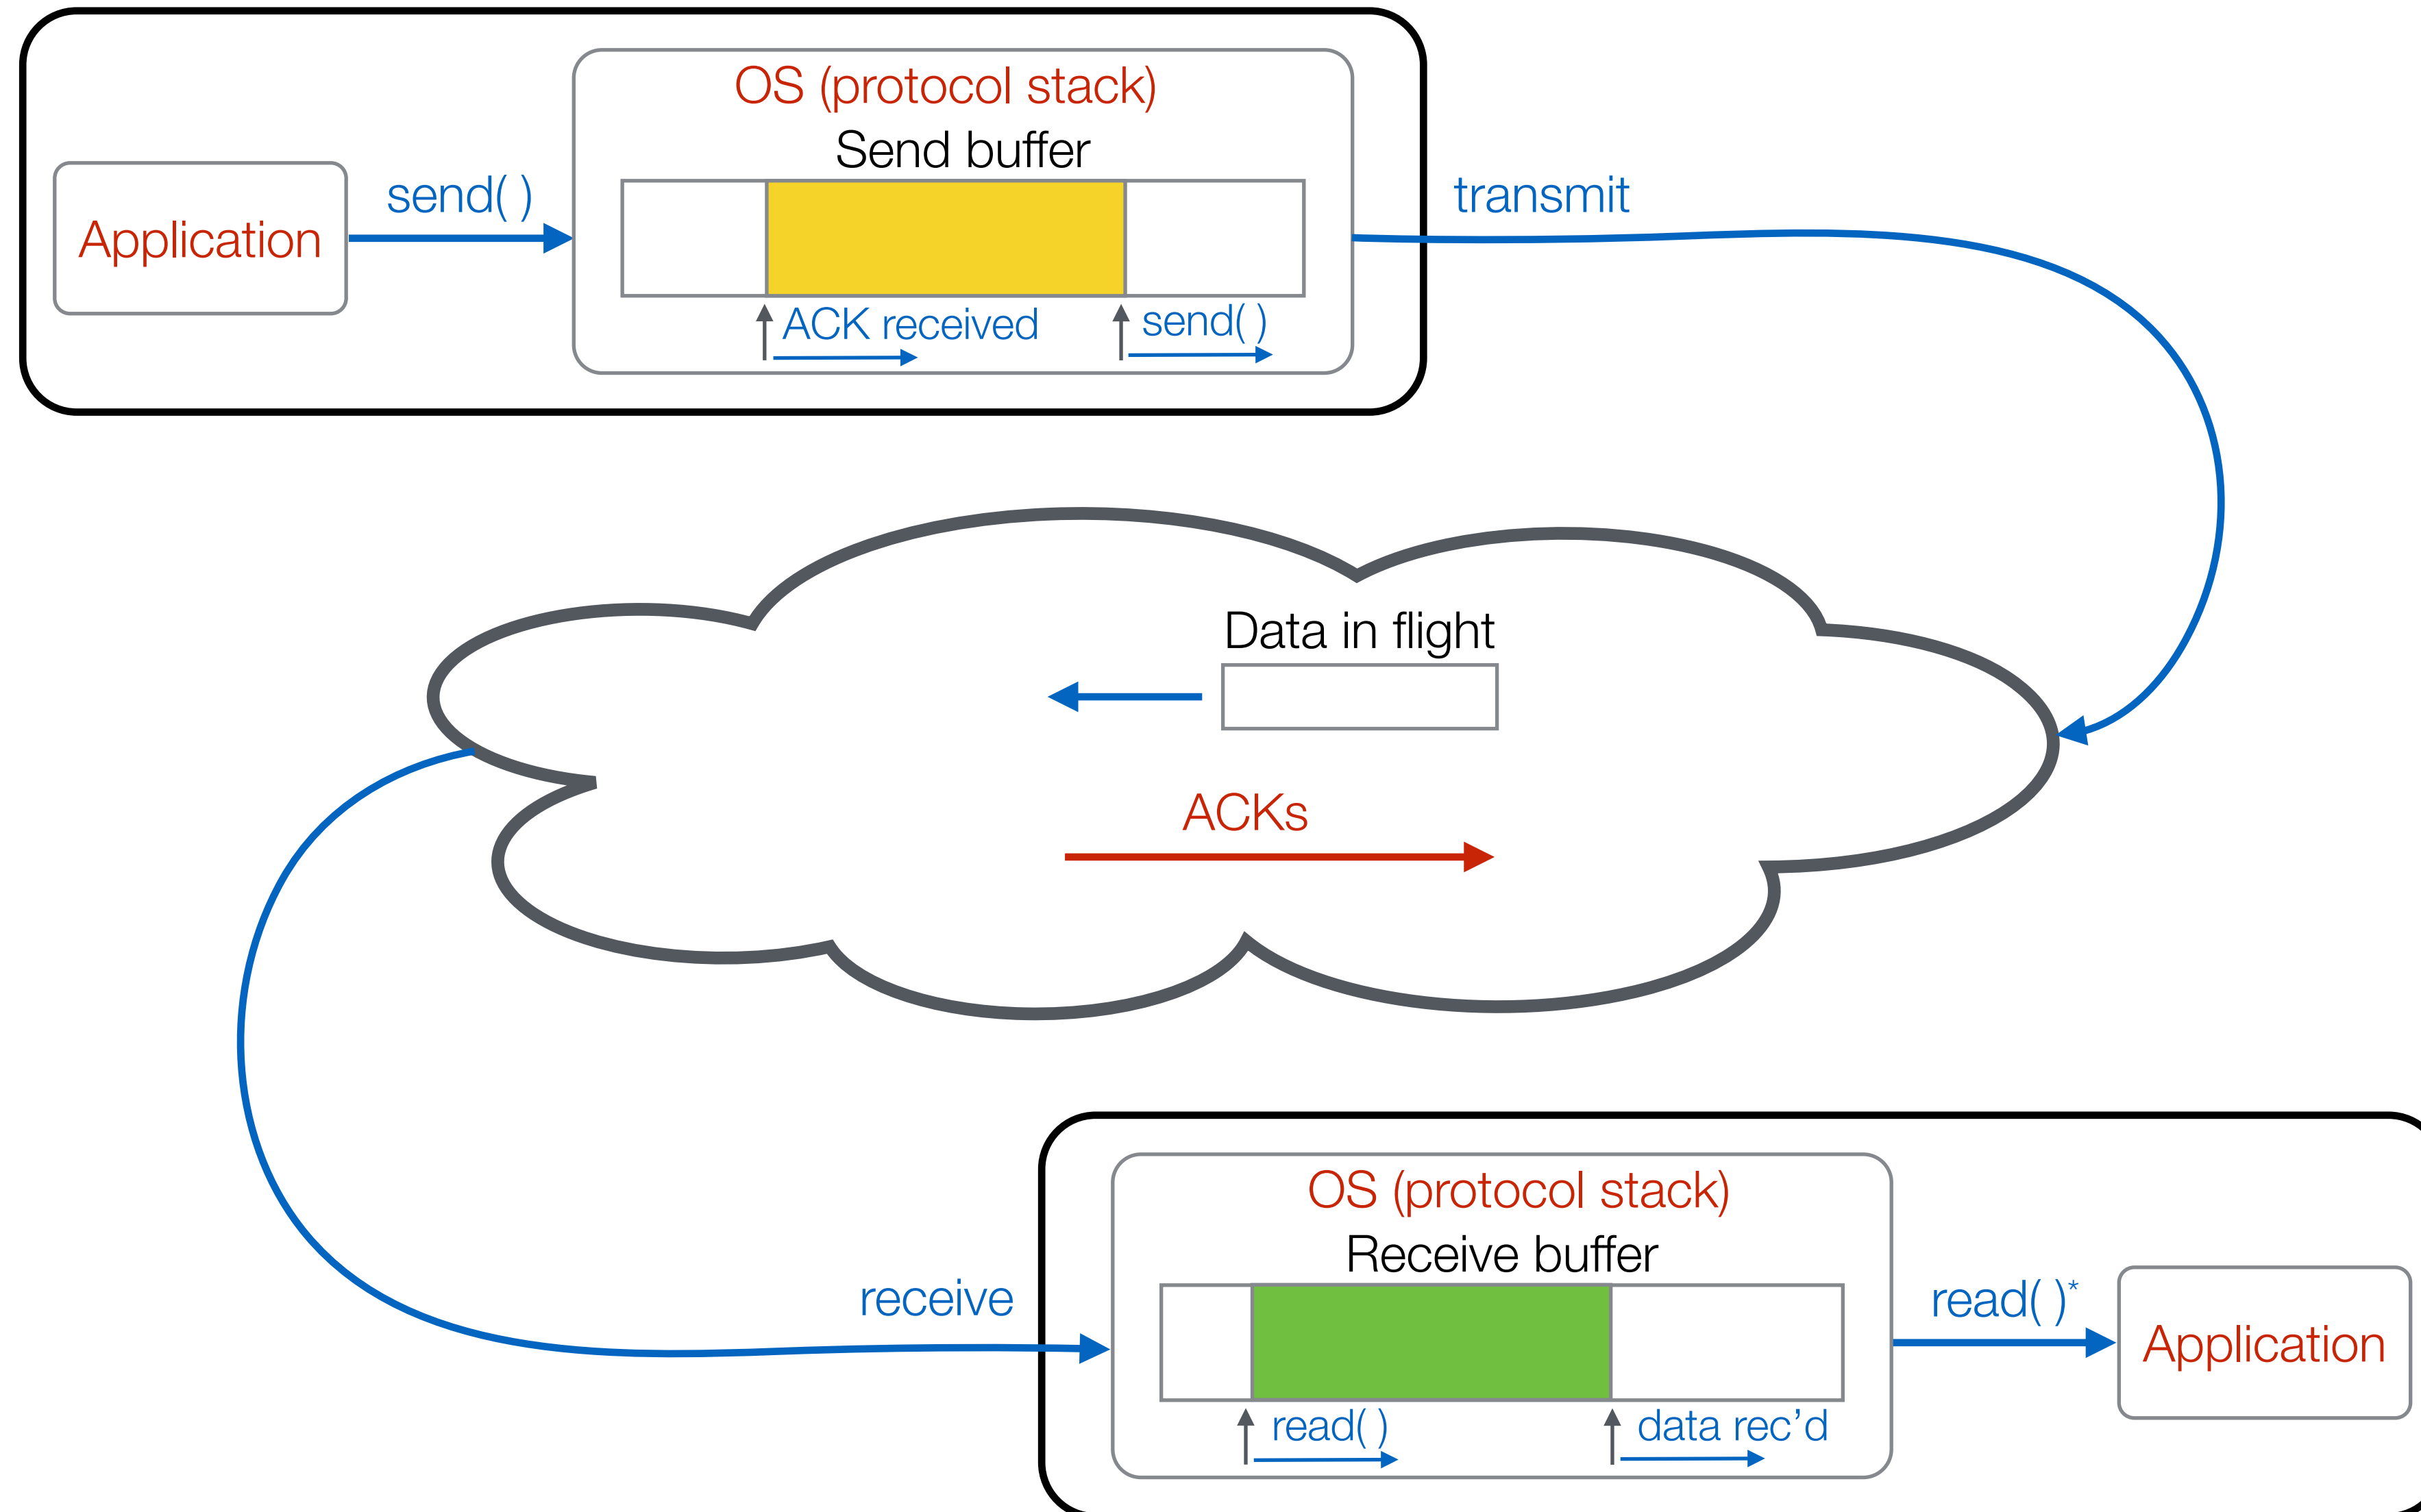
- ▶ Transport Control Protocol
- ▶ Design parameters and objectives
 - used by most popular applications, majority of Internet traffic is transported over TCP
 - significant impact on congestion behavior of the Internet
 - must operate over networks with widely-varying characteristics
 - must be robust and (relatively) simple to implement

TCP Header

		TCP Header																															
Offsets	Octet	0								1								2								3							
Octet	Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	Source port																Destination port															
4	32	Sequence number																															
8	64	Acknowledgment number (if ACK set)																															
12	96	Data offset	Reserved	N	C	E	U	A	P	R	S	F	Window Size																				
			0 0 0	S	W	C	R	C	S	S	Y	I																					
					R	E	G	K	H	T	N	N																					
16	128	Checksum																Urgent pointer (if URG set)															
20	160	Options (if Data Offset > 5, padded at the end with "0" bytes if necessary)																															
...																															

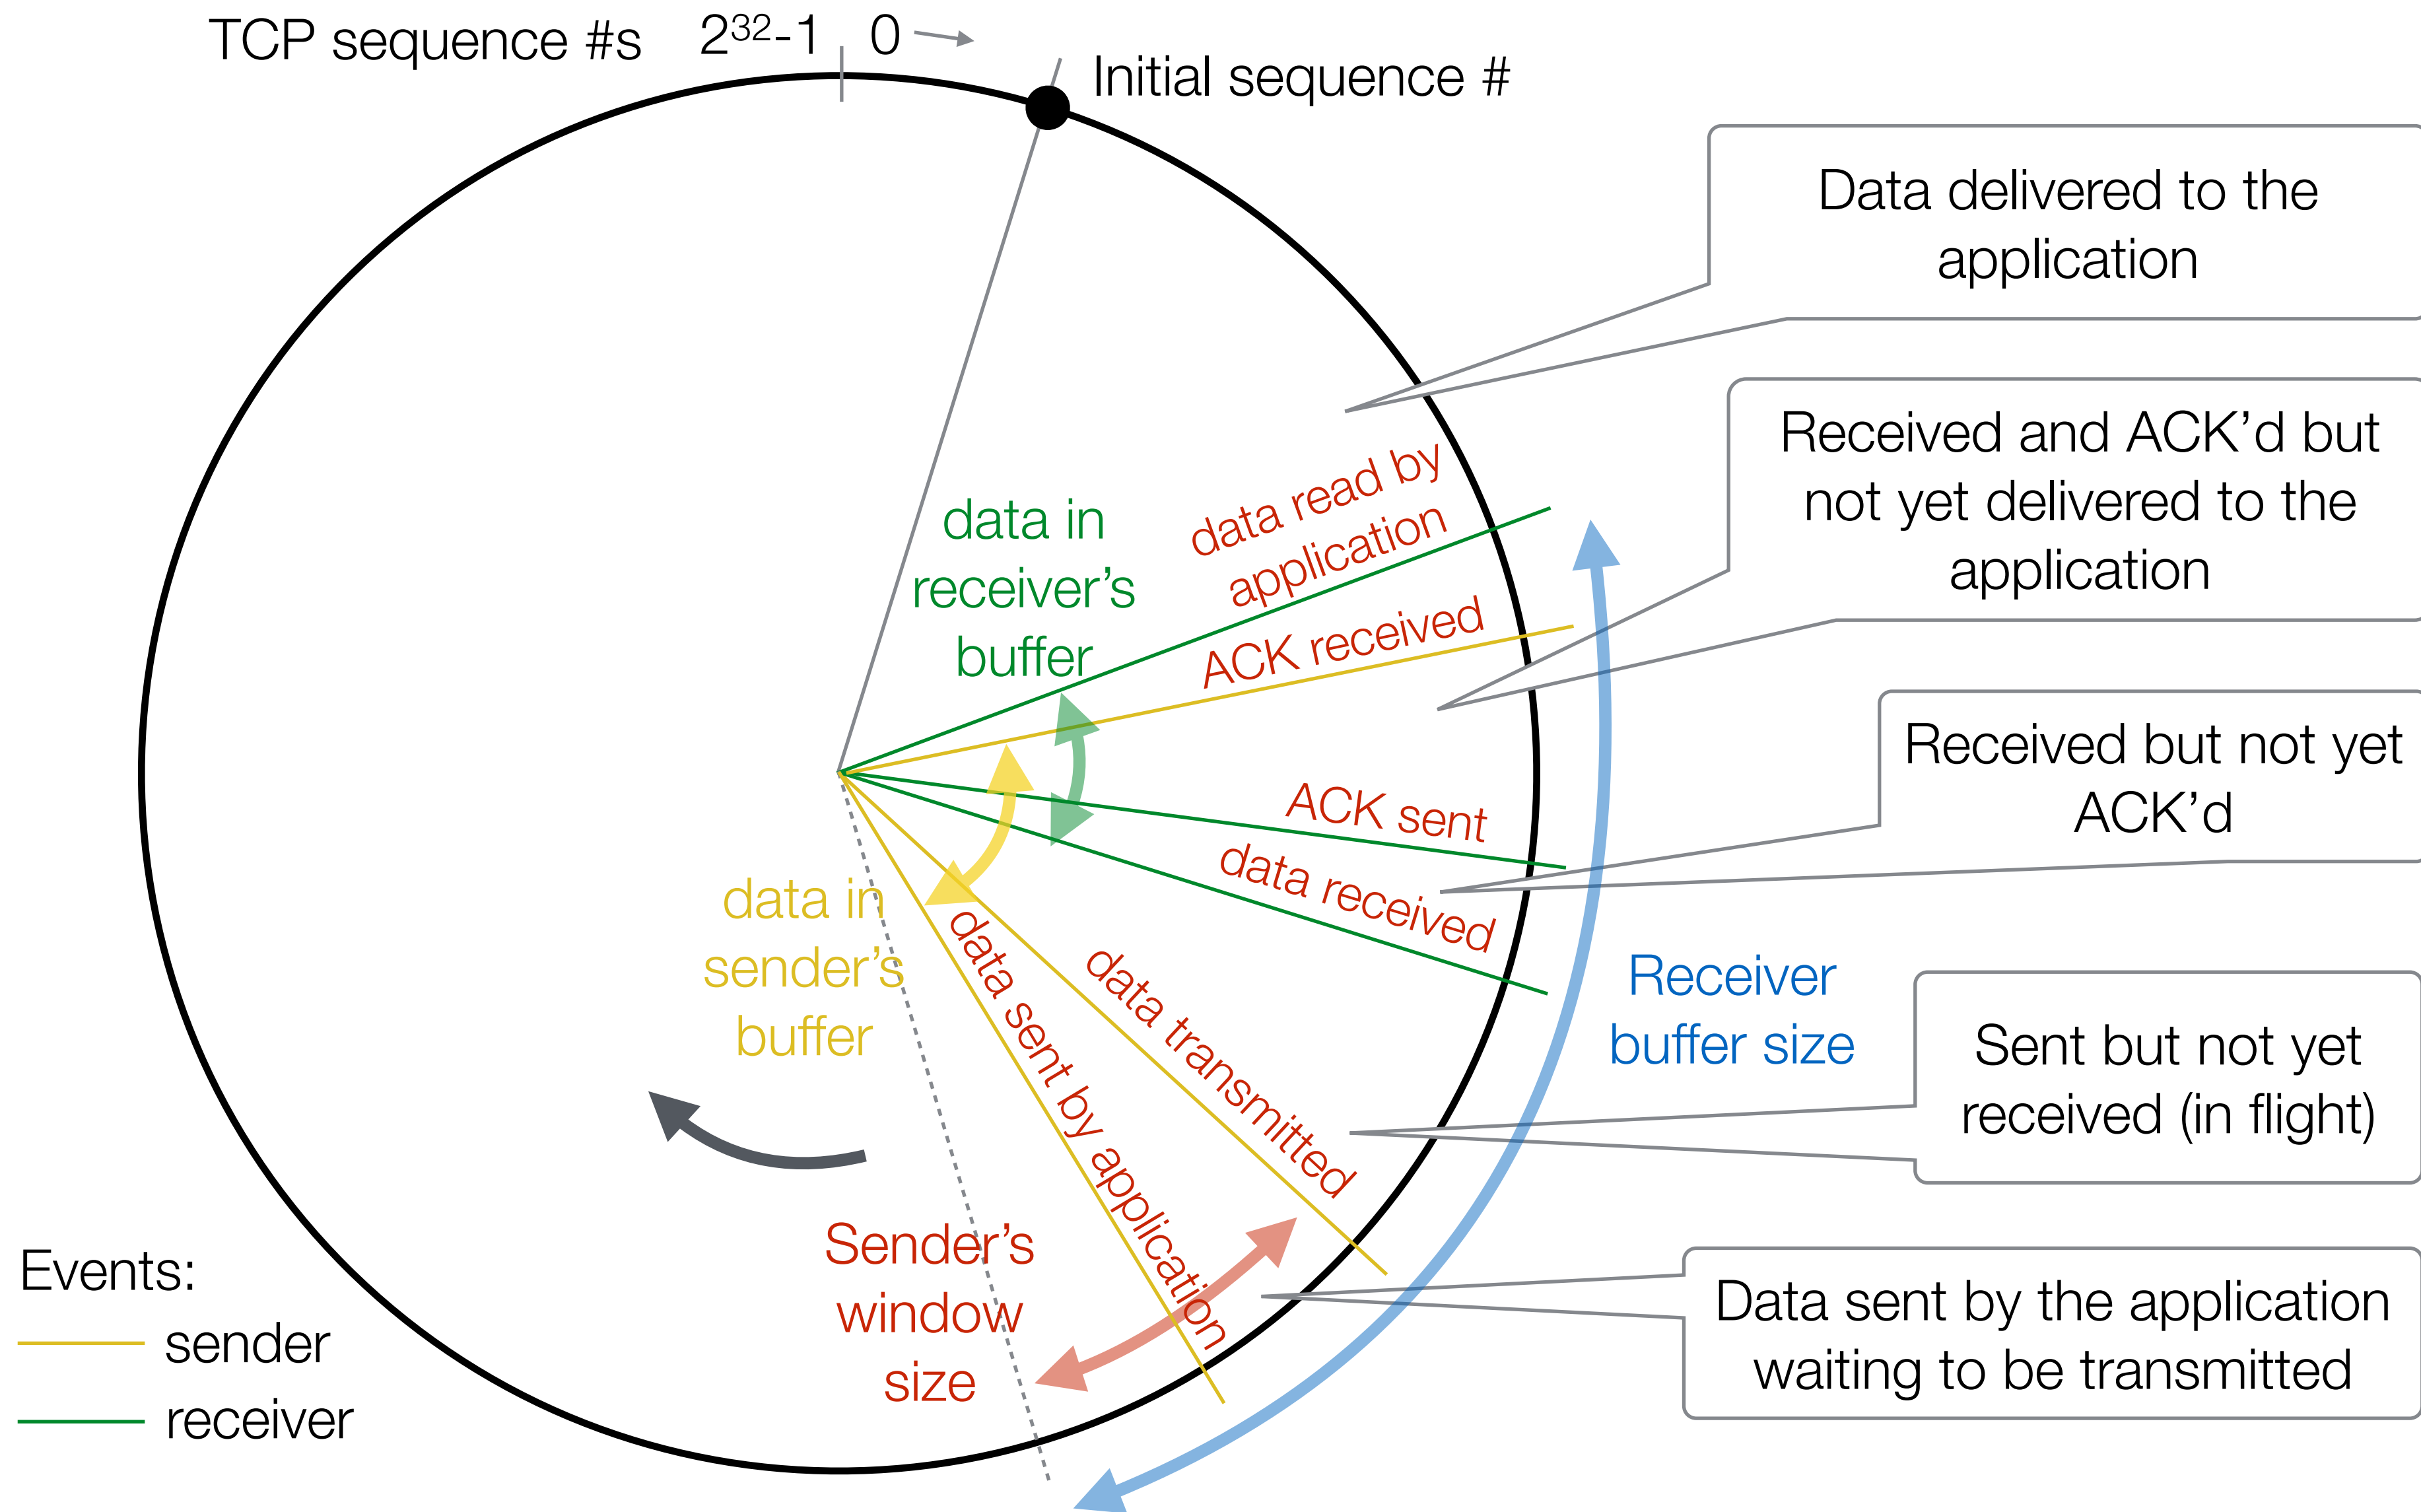
Another image appropriated from Wikipedia...

TCP buffering and data flow

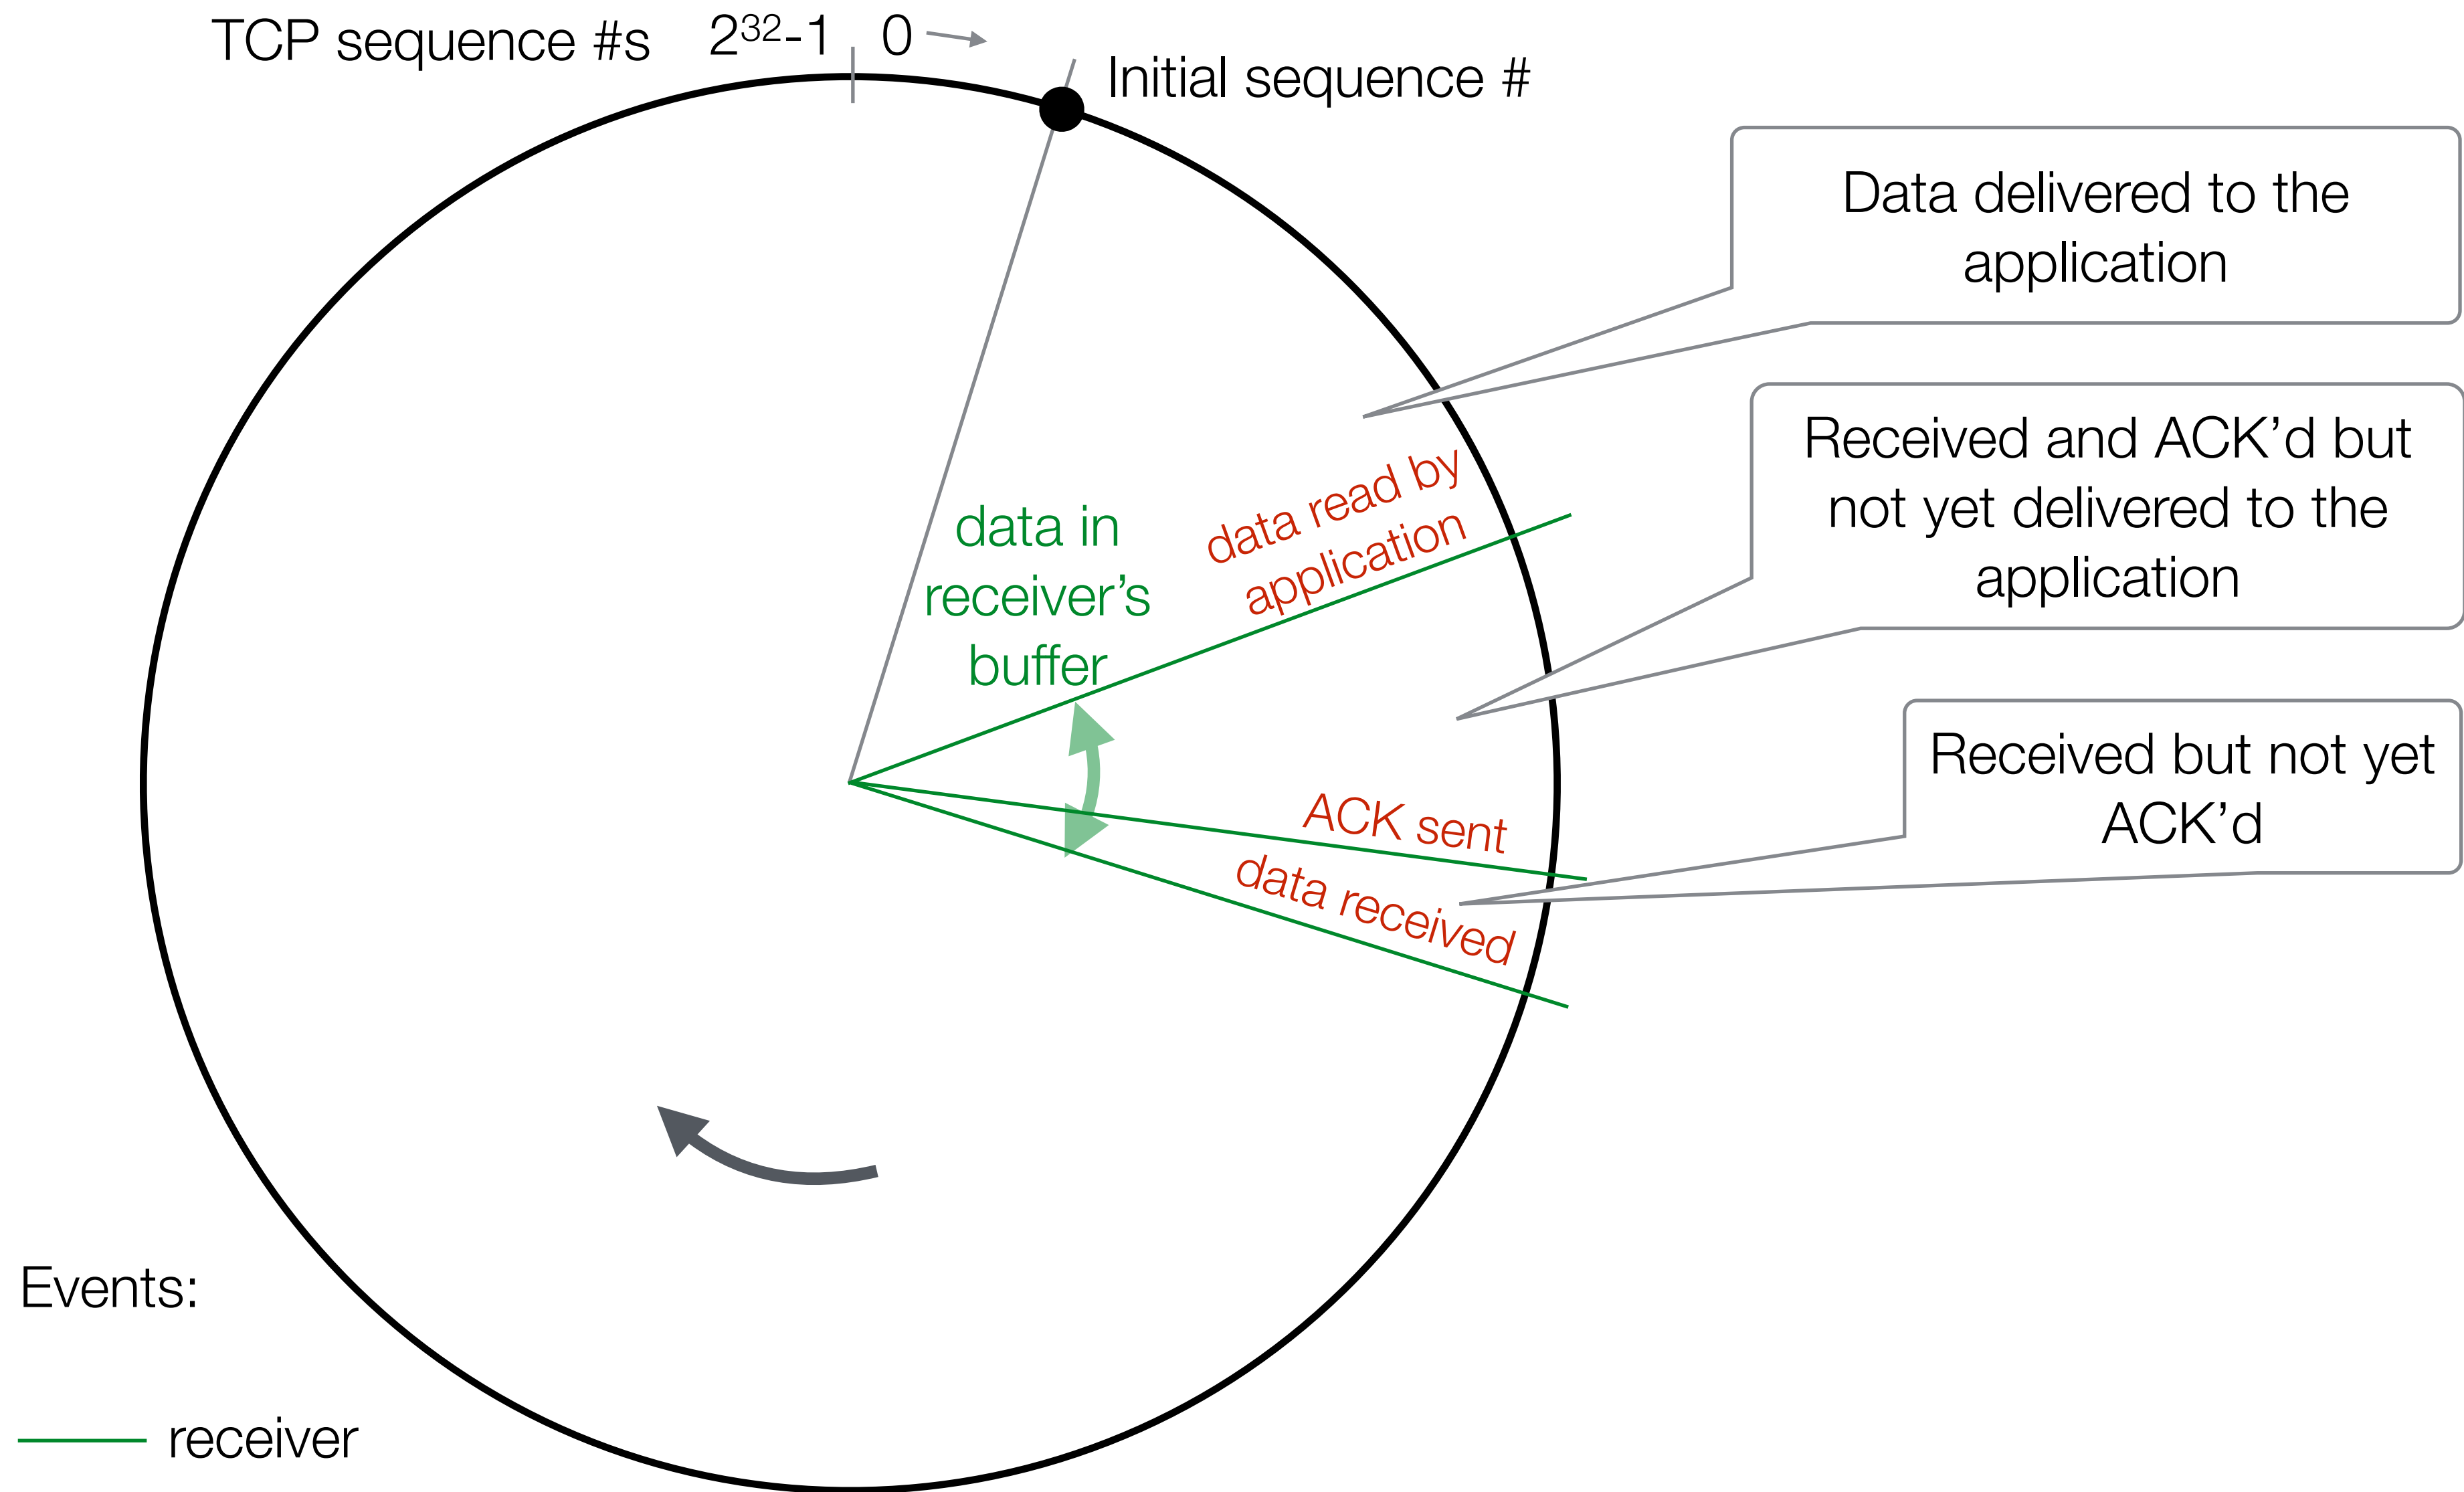


(*) many APIs call the `read()` operation "receive" (eg: `recv()`), `read` is used here to avoid confusion with receiving data on an interface

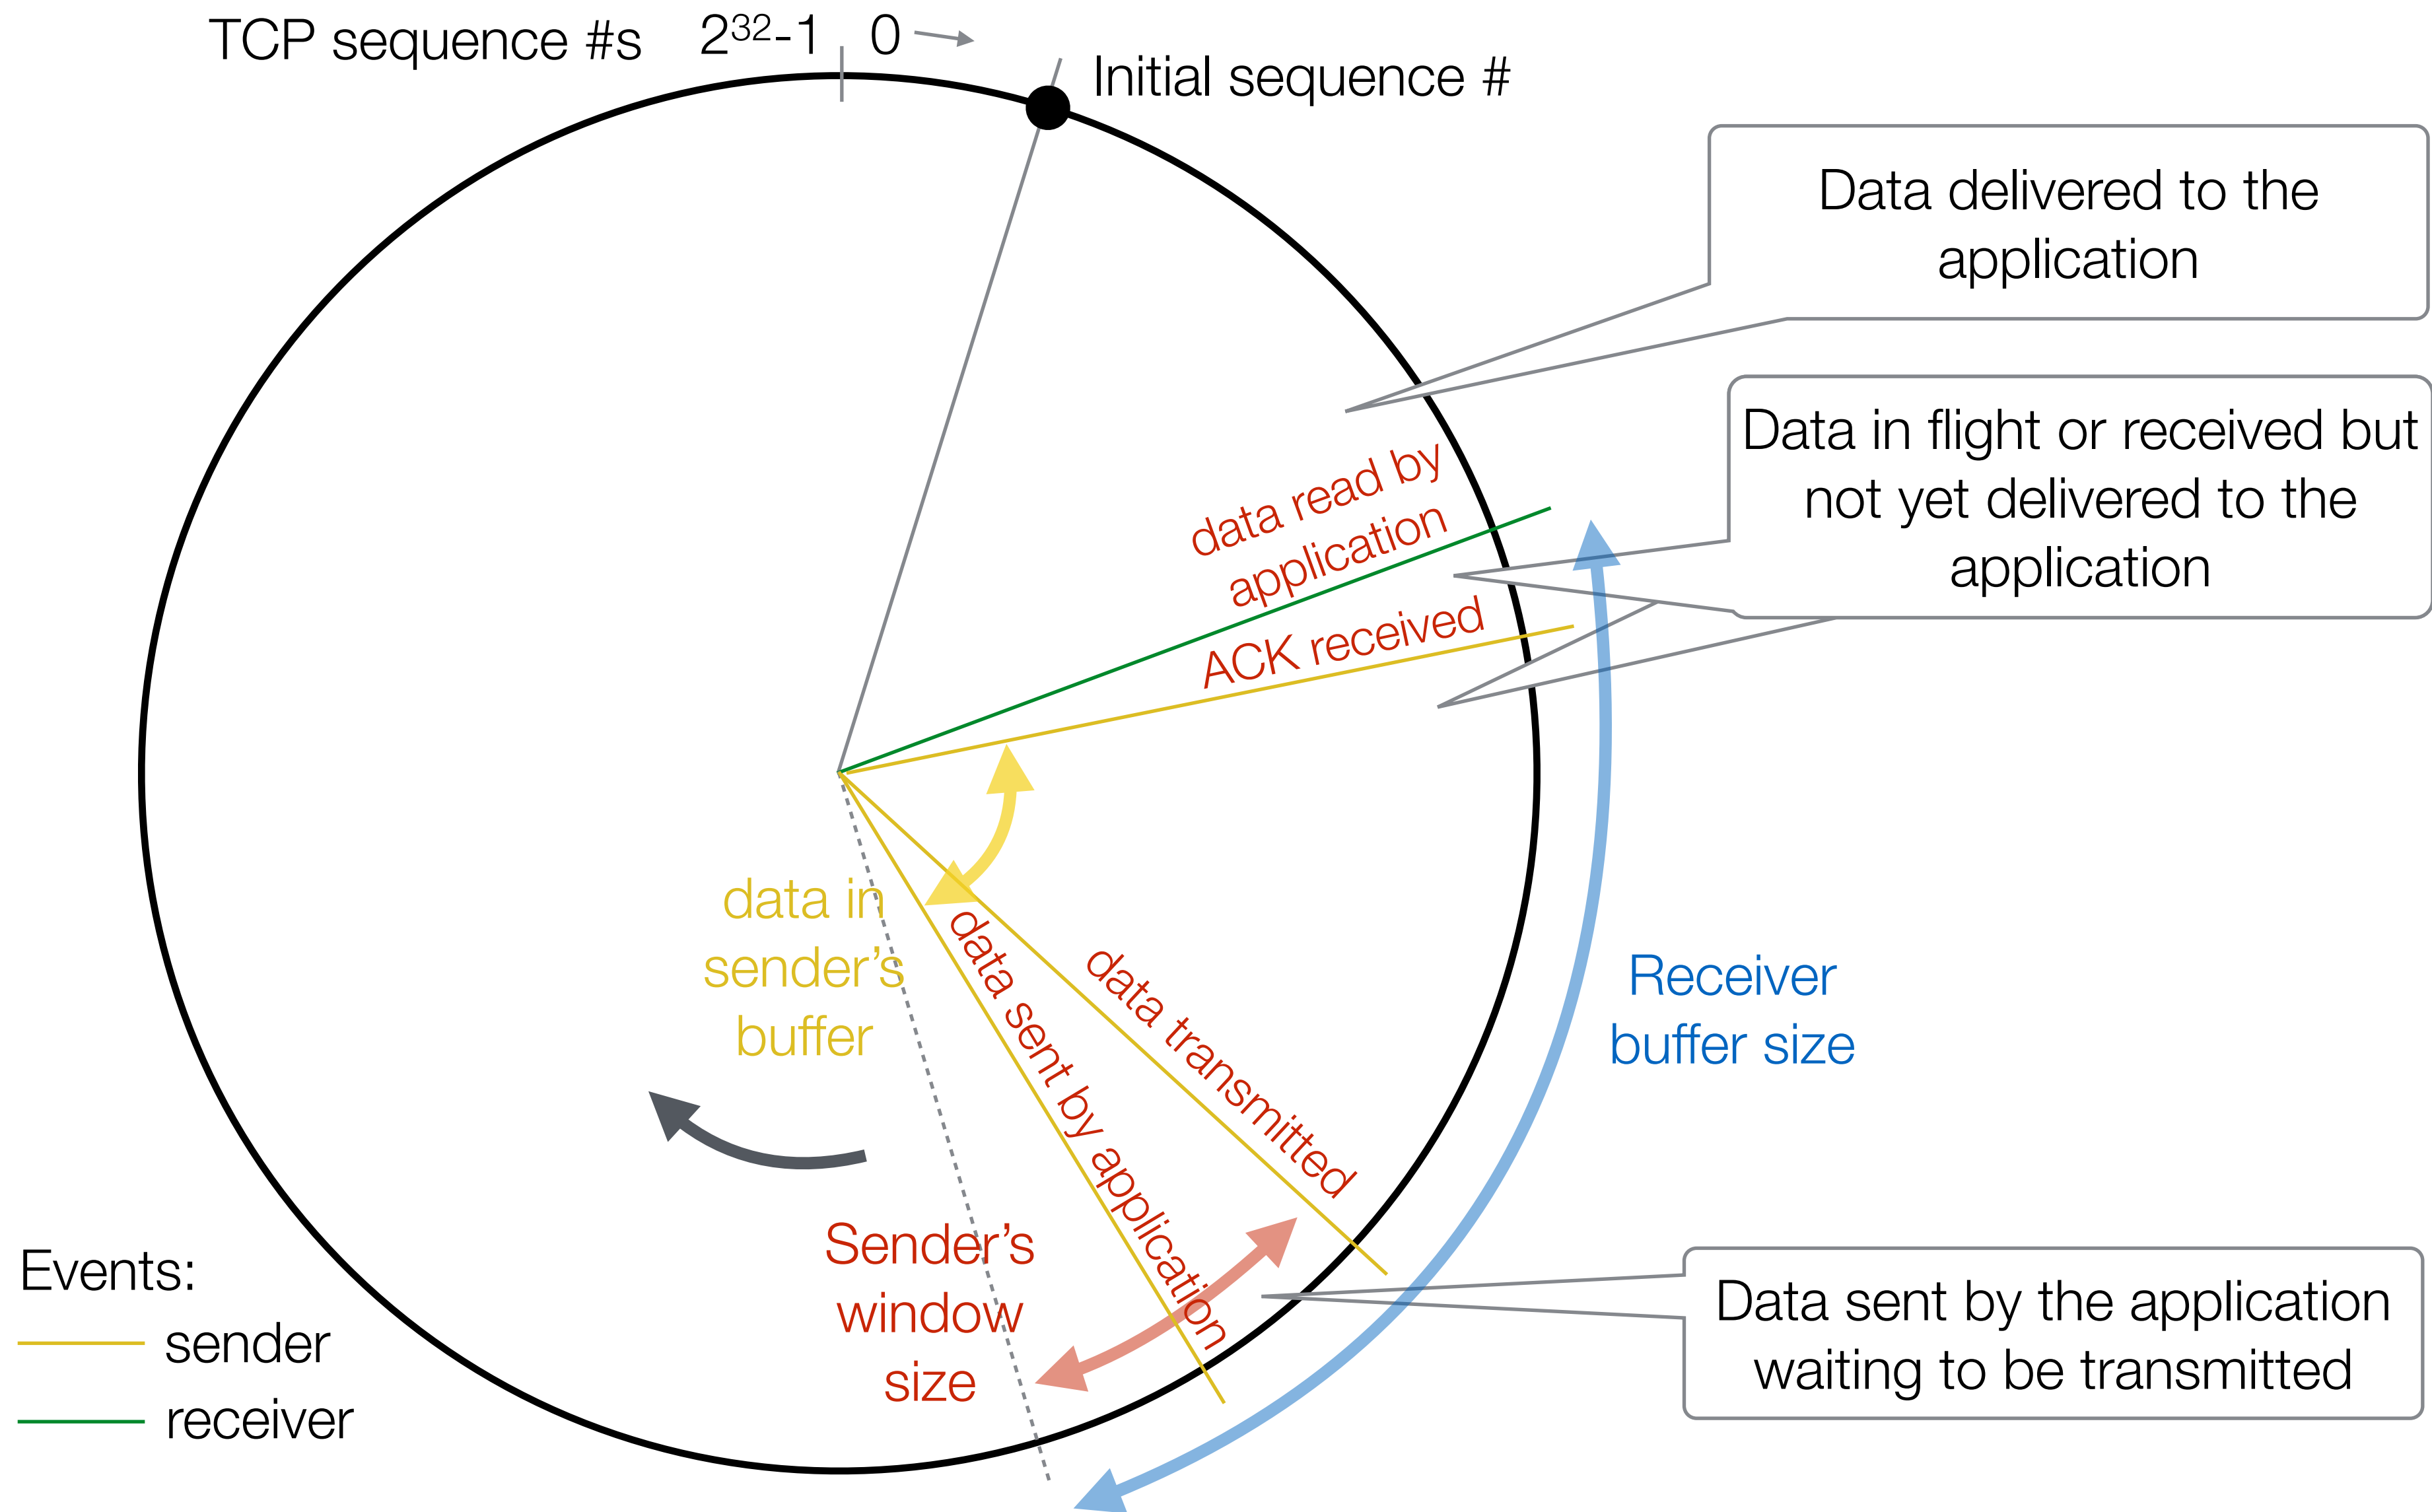
TCP Sliding Window



TCP Sliding Window



TCP Sliding Window



TCP session management

Offsets Octet		TCP Header																															
Octet	Bit	0								1								2								3							
Octet	Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	Source port																Destination port															
4	32	Sequence number																															
8	64	Acknowledgment number (if ACK set)																															
12	96	Data offset	Reserved 0 0 0			N S	C W R	E C R	U R G	A C K	P S H	R S T	S S Y	F I N	Window Size																		
16	128	Checksum																Urgent pointer (if URG set)															
20	160	Options (if Data Offset > 5, padded at the end with "0" bytes if necessary)																															
...																															

TCP Flags

Another image appropriated from Wikipedia...