

TCP Reno

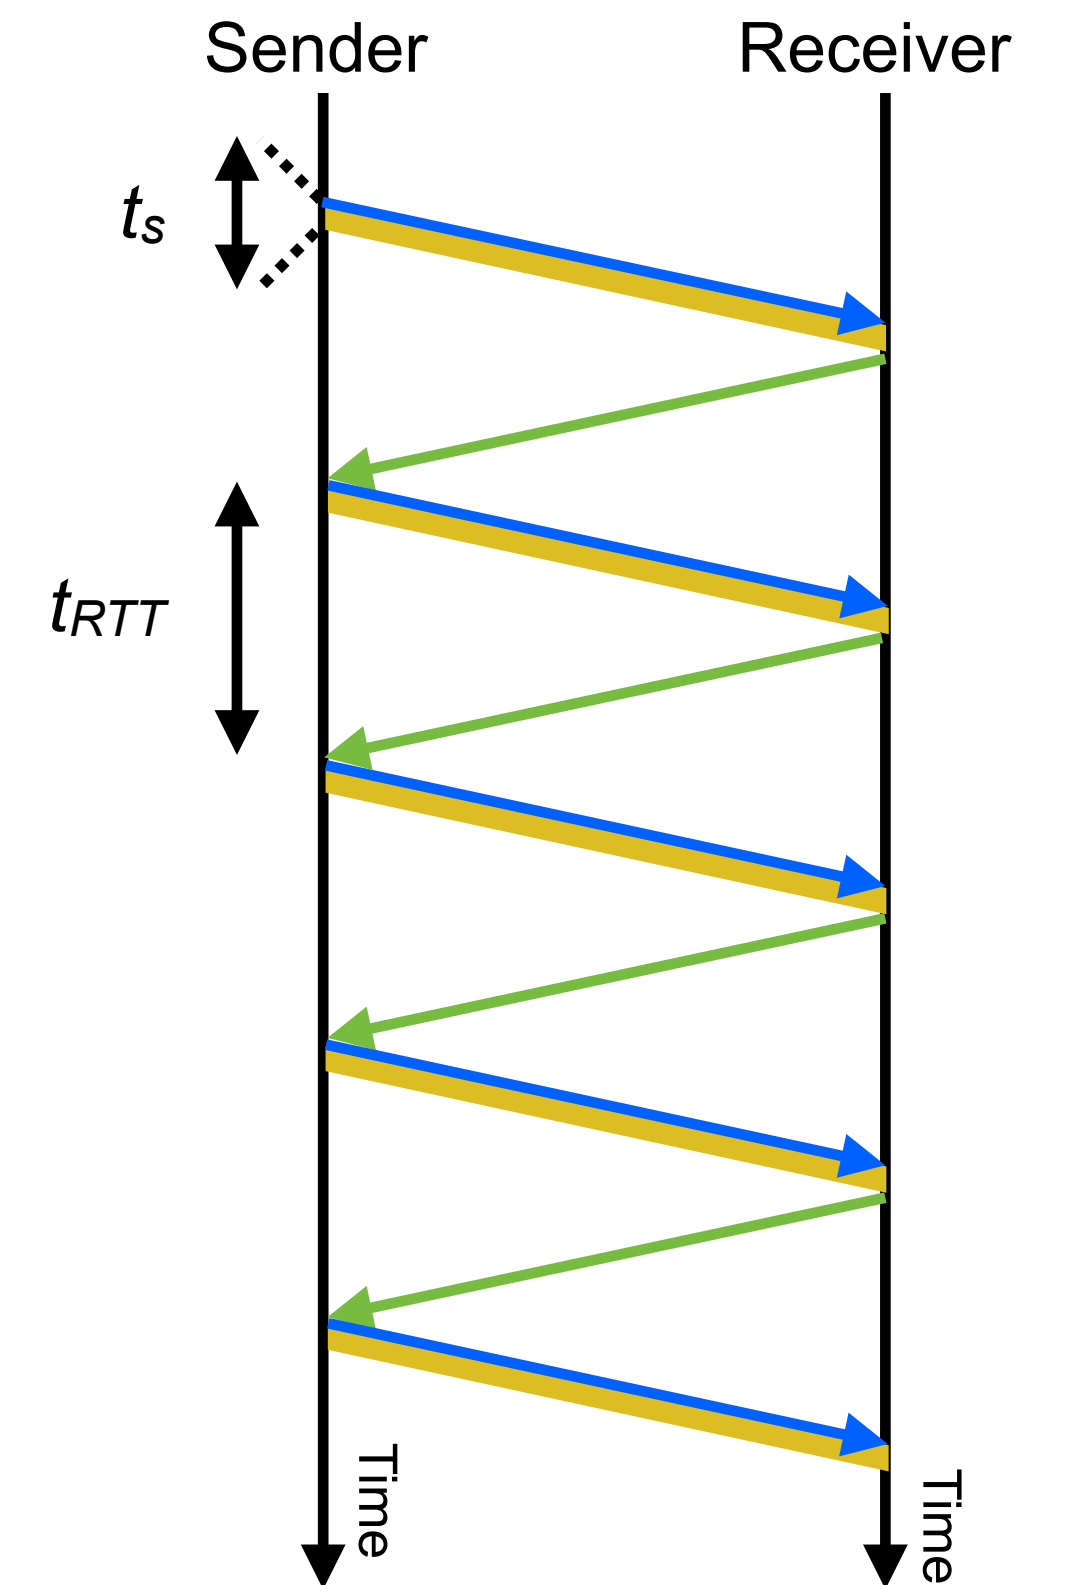
- Basic TCP congestion control
- Sliding window algorithms
- Slow start
- Congestion avoidance

Congestion Control in TCP Reno

- TCP uses window-based congestion control
 - Maintains a **sliding window** onto the available data that determines how much can be sent according to the AIMD algorithm
 - Plus slow start and congestion avoidance
 - Gives approximately equal share of the bandwidth to each flow sharing a link
- Basic congestion control algorithm known as **TCP Reno**
 - The state of the art in TCP as of ~1990


From Stop-and-Wait to Sliding Window Protocols

- Stop-and wait protocols perform poorly
 - It takes time, t_s , to send a packet
 - $t_s = (\text{packet size}) / (\text{link bandwidth})$
 - Acknowledgement returns t_{RTT} seconds later
 - Link utilisation, $U = t_s / t_{RTT}$
 - Fraction of time link is sending packets \rightarrow want $U \approx 1.0$
 - Assume a gigabit link sending a 1500 byte packet from Glasgow to London:
 - $t_s = 1500 \times 8 \text{ bits} / 10^9 \text{ bits per second} = 0.000012\text{s}$
 - $t_{RTT} \approx 0.010 \text{ seconds}$
 - $U \approx 0.0012$
 - *i.e.*, the link is in use 0.12% of the time
- Sliding window protocols improve on stop-and-wait by sending more than one packet before stopping for acknowledgement



Sliding Window Protocols Improve Link Utilisation


Sliding window protocols improve link utilisation using a **congestion window** – number of packets to be sent before an acknowledgement arrives



In this example, the window size is six packets → acknowledgement for packet 6 arrives just in time to release packet 7 for transmission


Each returning acknowledgement for new data slides the window along, releasing next packet for transmission → if window sized correctly, each acknowledgement **arrives just in time** to release next packet

What is the optimal size for the window? **bandwidth × delay** of path → **but neither is known to the sender**




TCP Reno Congestion Control

- **TCP Reno is a sliding window protocol**
 - Optimised for not having information to know the correct window size

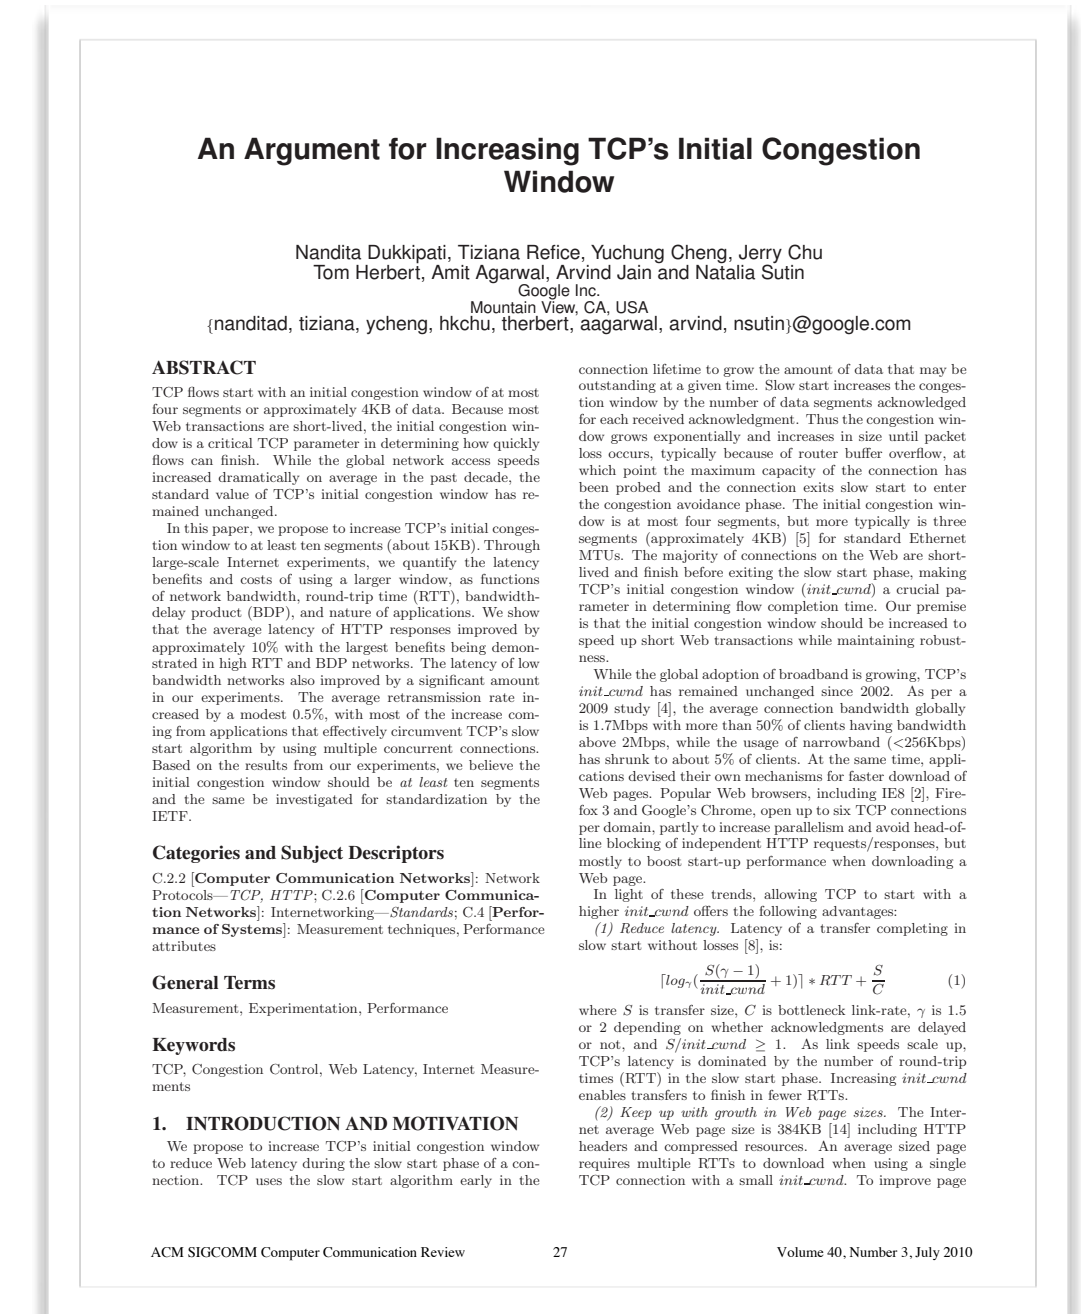


- How to choose initial window?
- How to find the path capacity?
 - Slow start to estimate the bottleneck link capacity
 - Congestion avoidance to probe for changes in capacity



TCP Reno: Choosing the Initial Window

- How to choose initial window size, W_{init} ?
 - No information → need to measure path capacity
 - Start with a small window, increase until congestion
 - **$W_{init} = 1$ packet per round-trip time (RTT) is safe**
 - Start at the slowest possible rate, equivalent to stop-and-wait, and increase
 - **$W_{init} = 3$ packets per RTT**
 - Traditional TCP Reno approach
 - **$W_{init} = 10$ packets per RTT**
 - Modern TCP implementations [RFC 6928]
 - Compromise between safety and performance – measurements show this is generally safe at present
 - Expect W_{init} to gradually be increased as average network performance improves



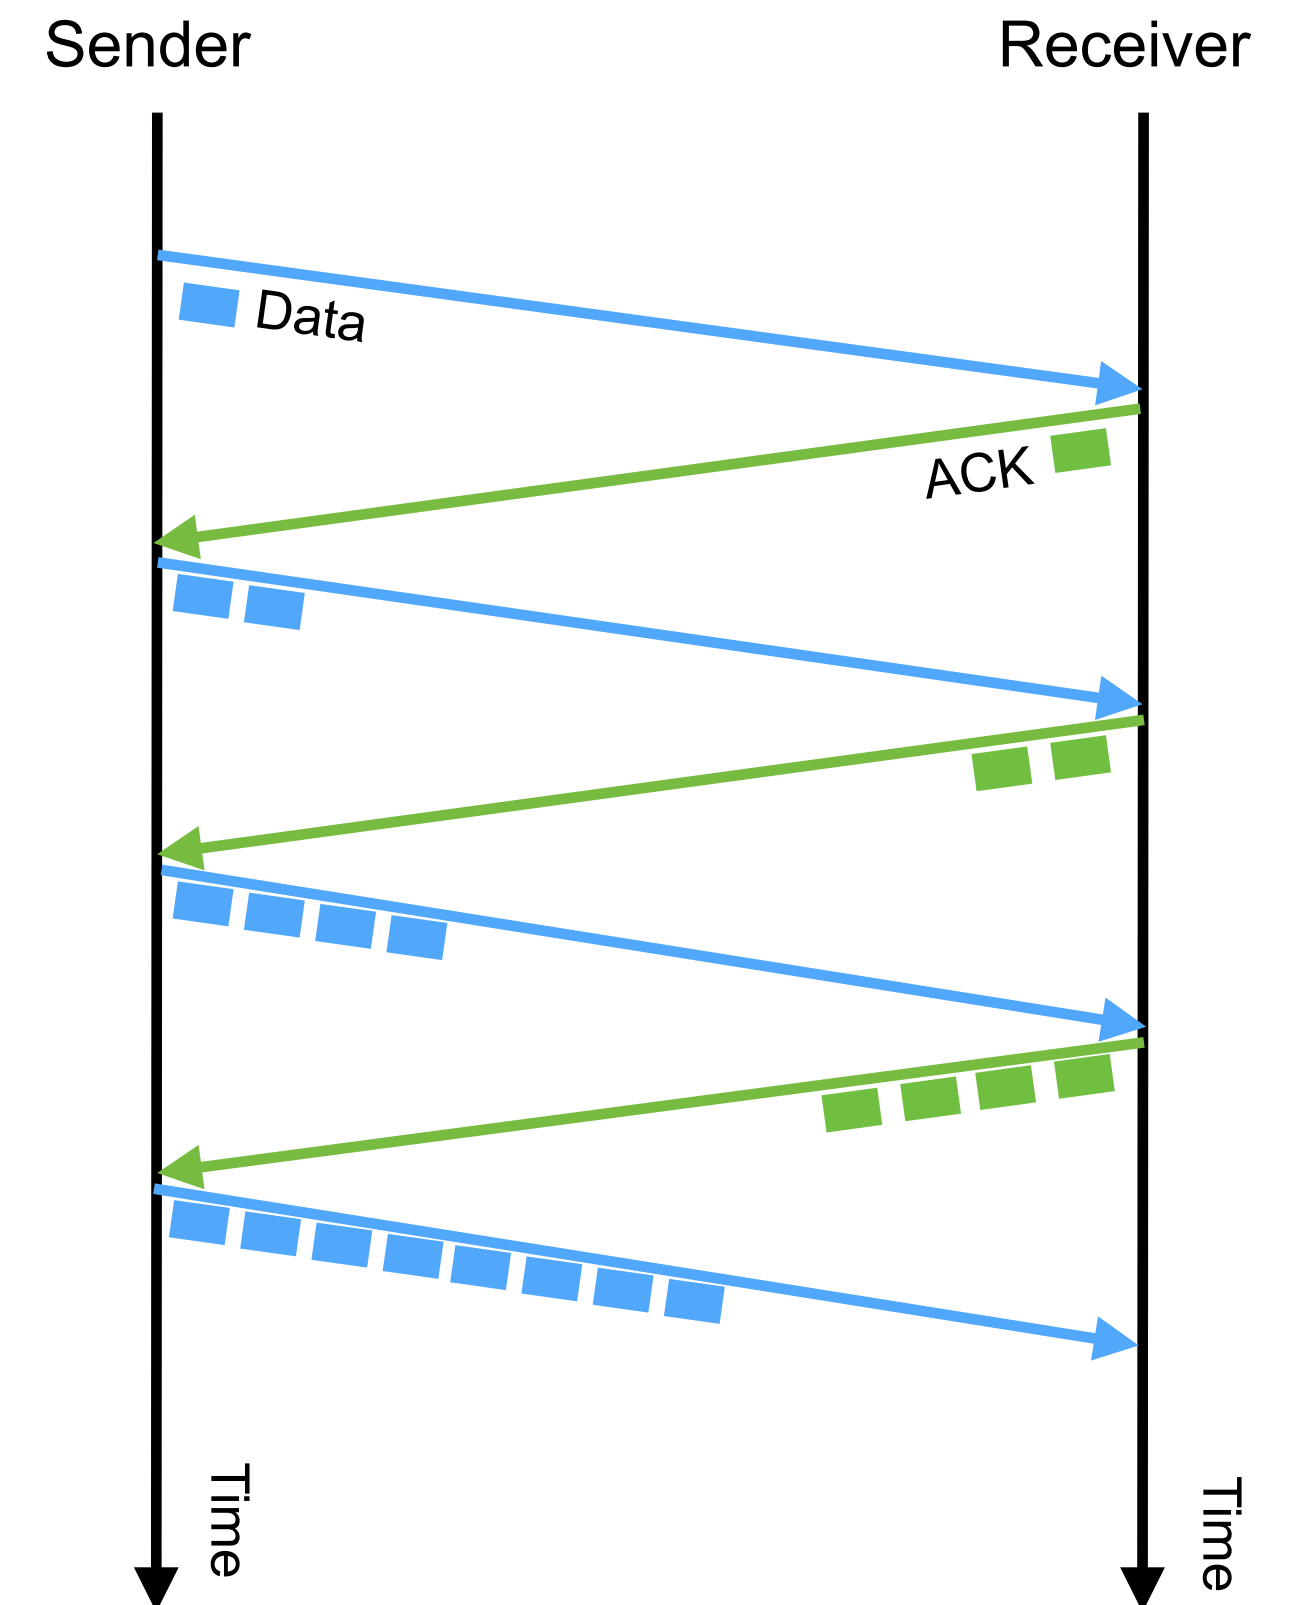
N. Dukkkipati *et al*, “An argument for increasing TCP’s initial congestion window”, ACM Computer Communication Review, 40(3):27–33, July 2010. <http://dx.doi.org/10.1145/1823844.1823848>

TCP Reno: Finding the Path Capacity

- The initial window allows you to send **something**
 - Unlikely to be the optimal window size
 - How to choose the correct window size to match the link capacity?
 - **Slow start** to rapidly find the correct window size for the path
 - **Congestion avoidance** to adapt to changes in path capacity once connection is running


TCP Reno: Slow Start (1/2)

- Assume $W_{init} = 1$ packet per RTT – **slow start to connection**
 - It will be $W_{init} = 3$ or $W_{init} = 10$ in practice, but using $W_{init} = 1$ makes the example simpler...
- Rapidly increase window until network capacity is reached
 - Each acknowledgement for new data increases congestion window, W , by 1 packet per RTT → congestion window doubles each RTT
 - If a packet is lost, halve congestion window back to previous value and exit slow start



TCP Reno: Slow Start (2/2)

- TCP Reno slow start phase:
 - Starts sending slowly
 - Rapidly increases sending rate until a packet is lost
 - Resets sending rate to last known good rate when first loss occurs

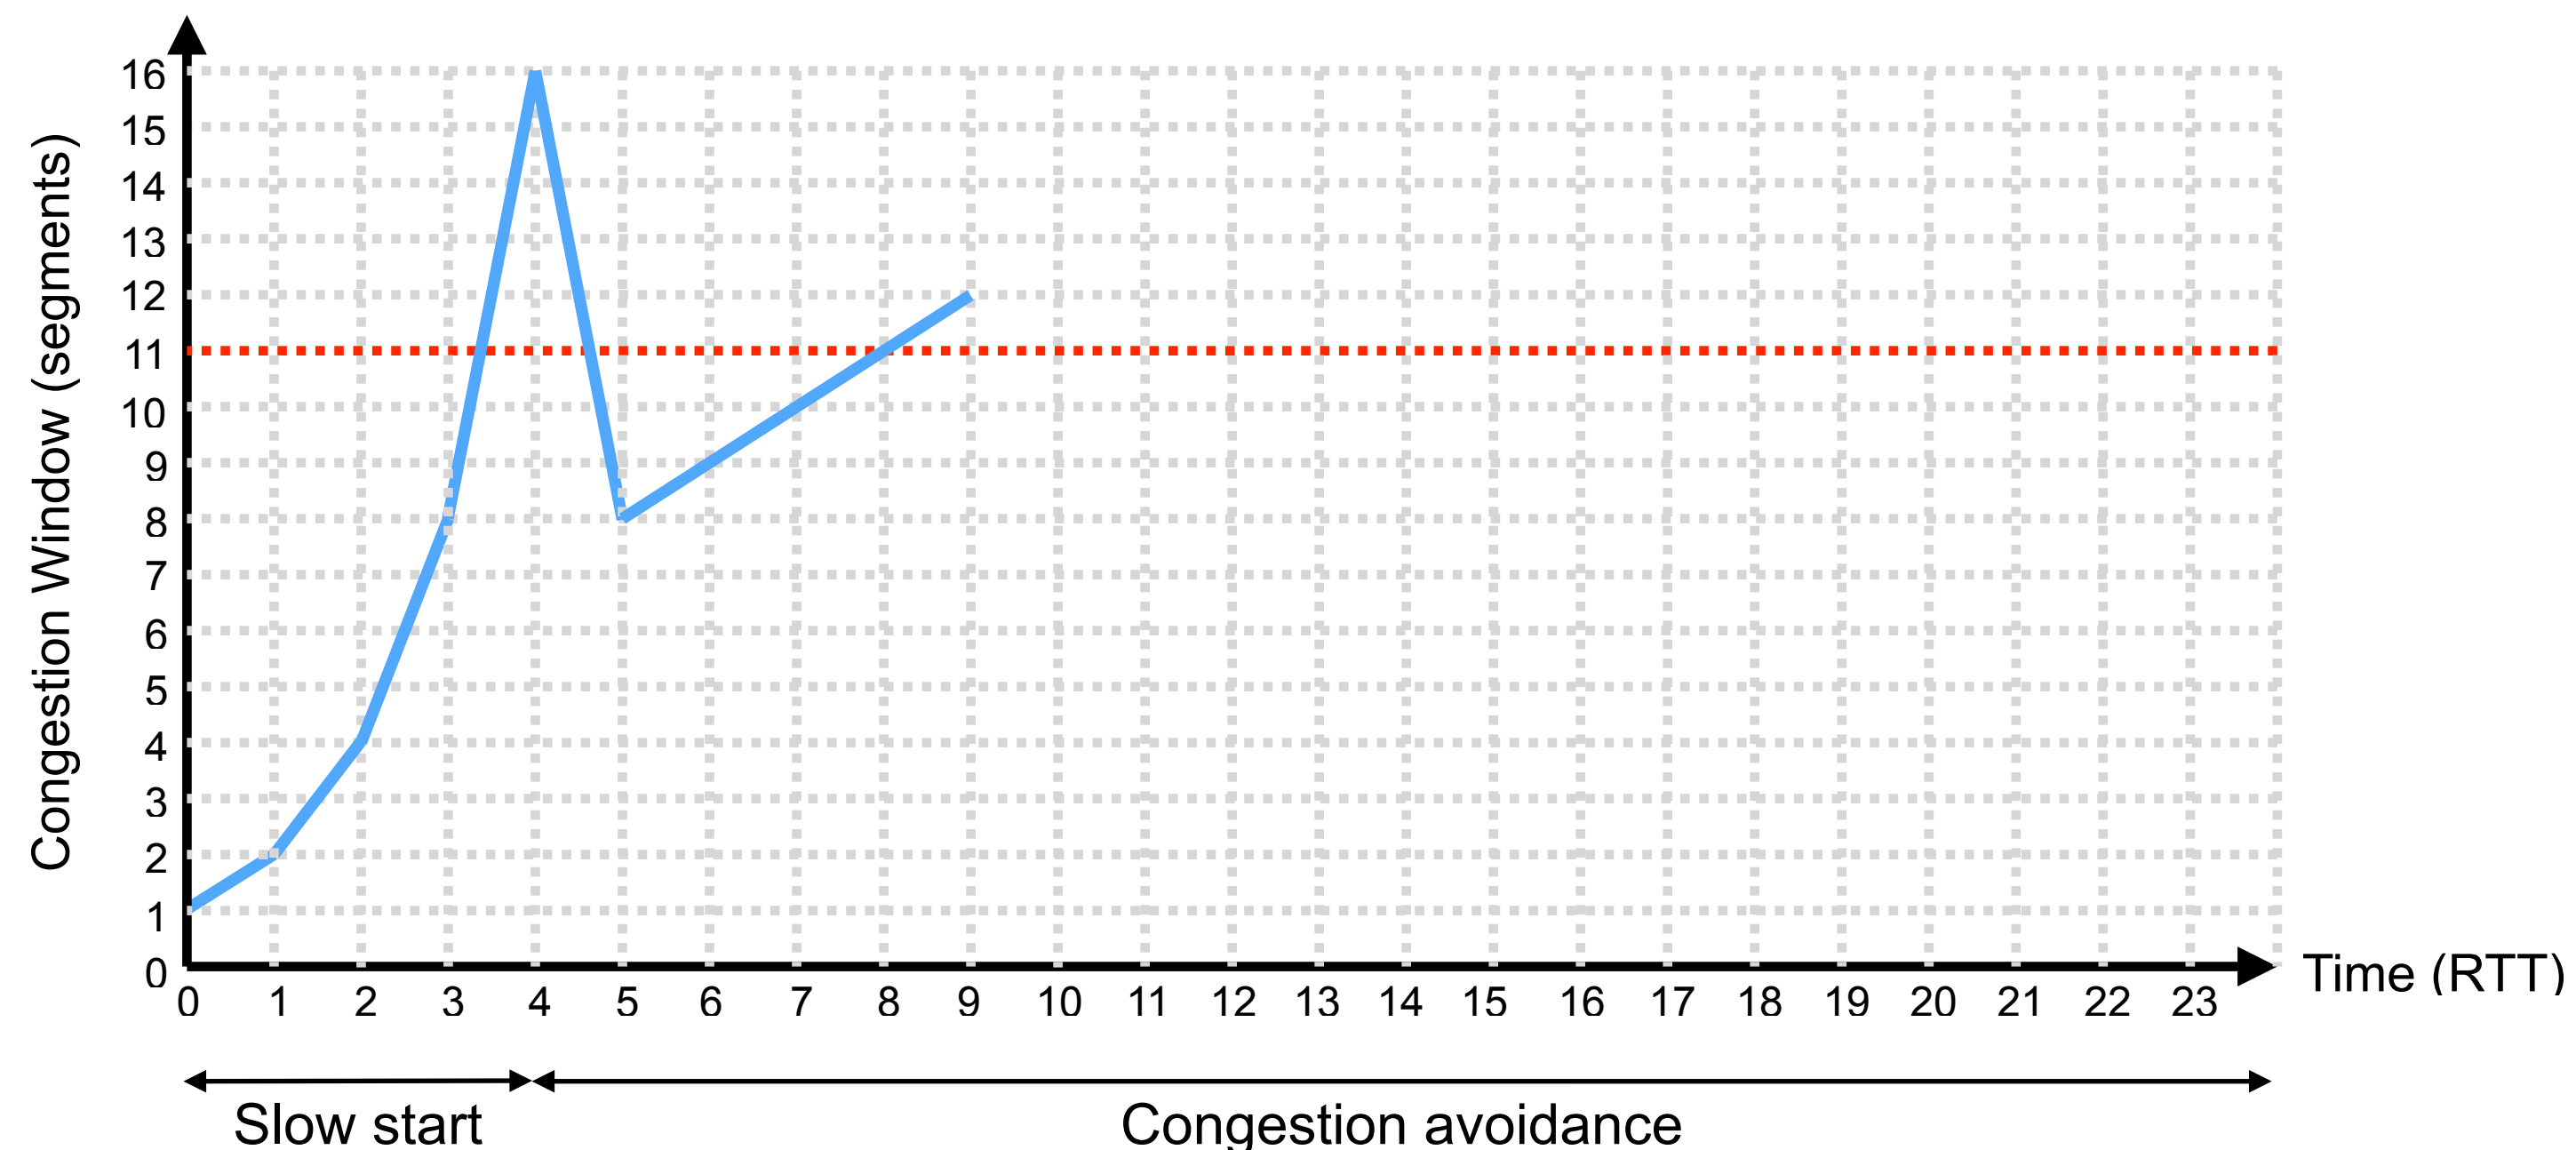


TCP Reno: Congestion Avoidance (1/4)

- After first packet is lost, TCP switches to **congestion avoidance**
 - The congestion window is now approximately the right size for the path
 - Goal is now to adapt to changes in network capacity
 - Perhaps the path capacity changes – radio signal strength changes for mobile device
 - Perhaps the other traffic changes – competing flows stop; additional cross-traffic starts
- Additive increase, multiplicative decrease of congestion window


TCP Reno: Congestion Avoidance (2/4)

- If a complete window of packets is sent without loss:
 - Increase congestion window by 1 packet per RTT, then send next window worth of packets
 - Slow, linear, additive increase in window: $W_i = W_{i-1} + 1$



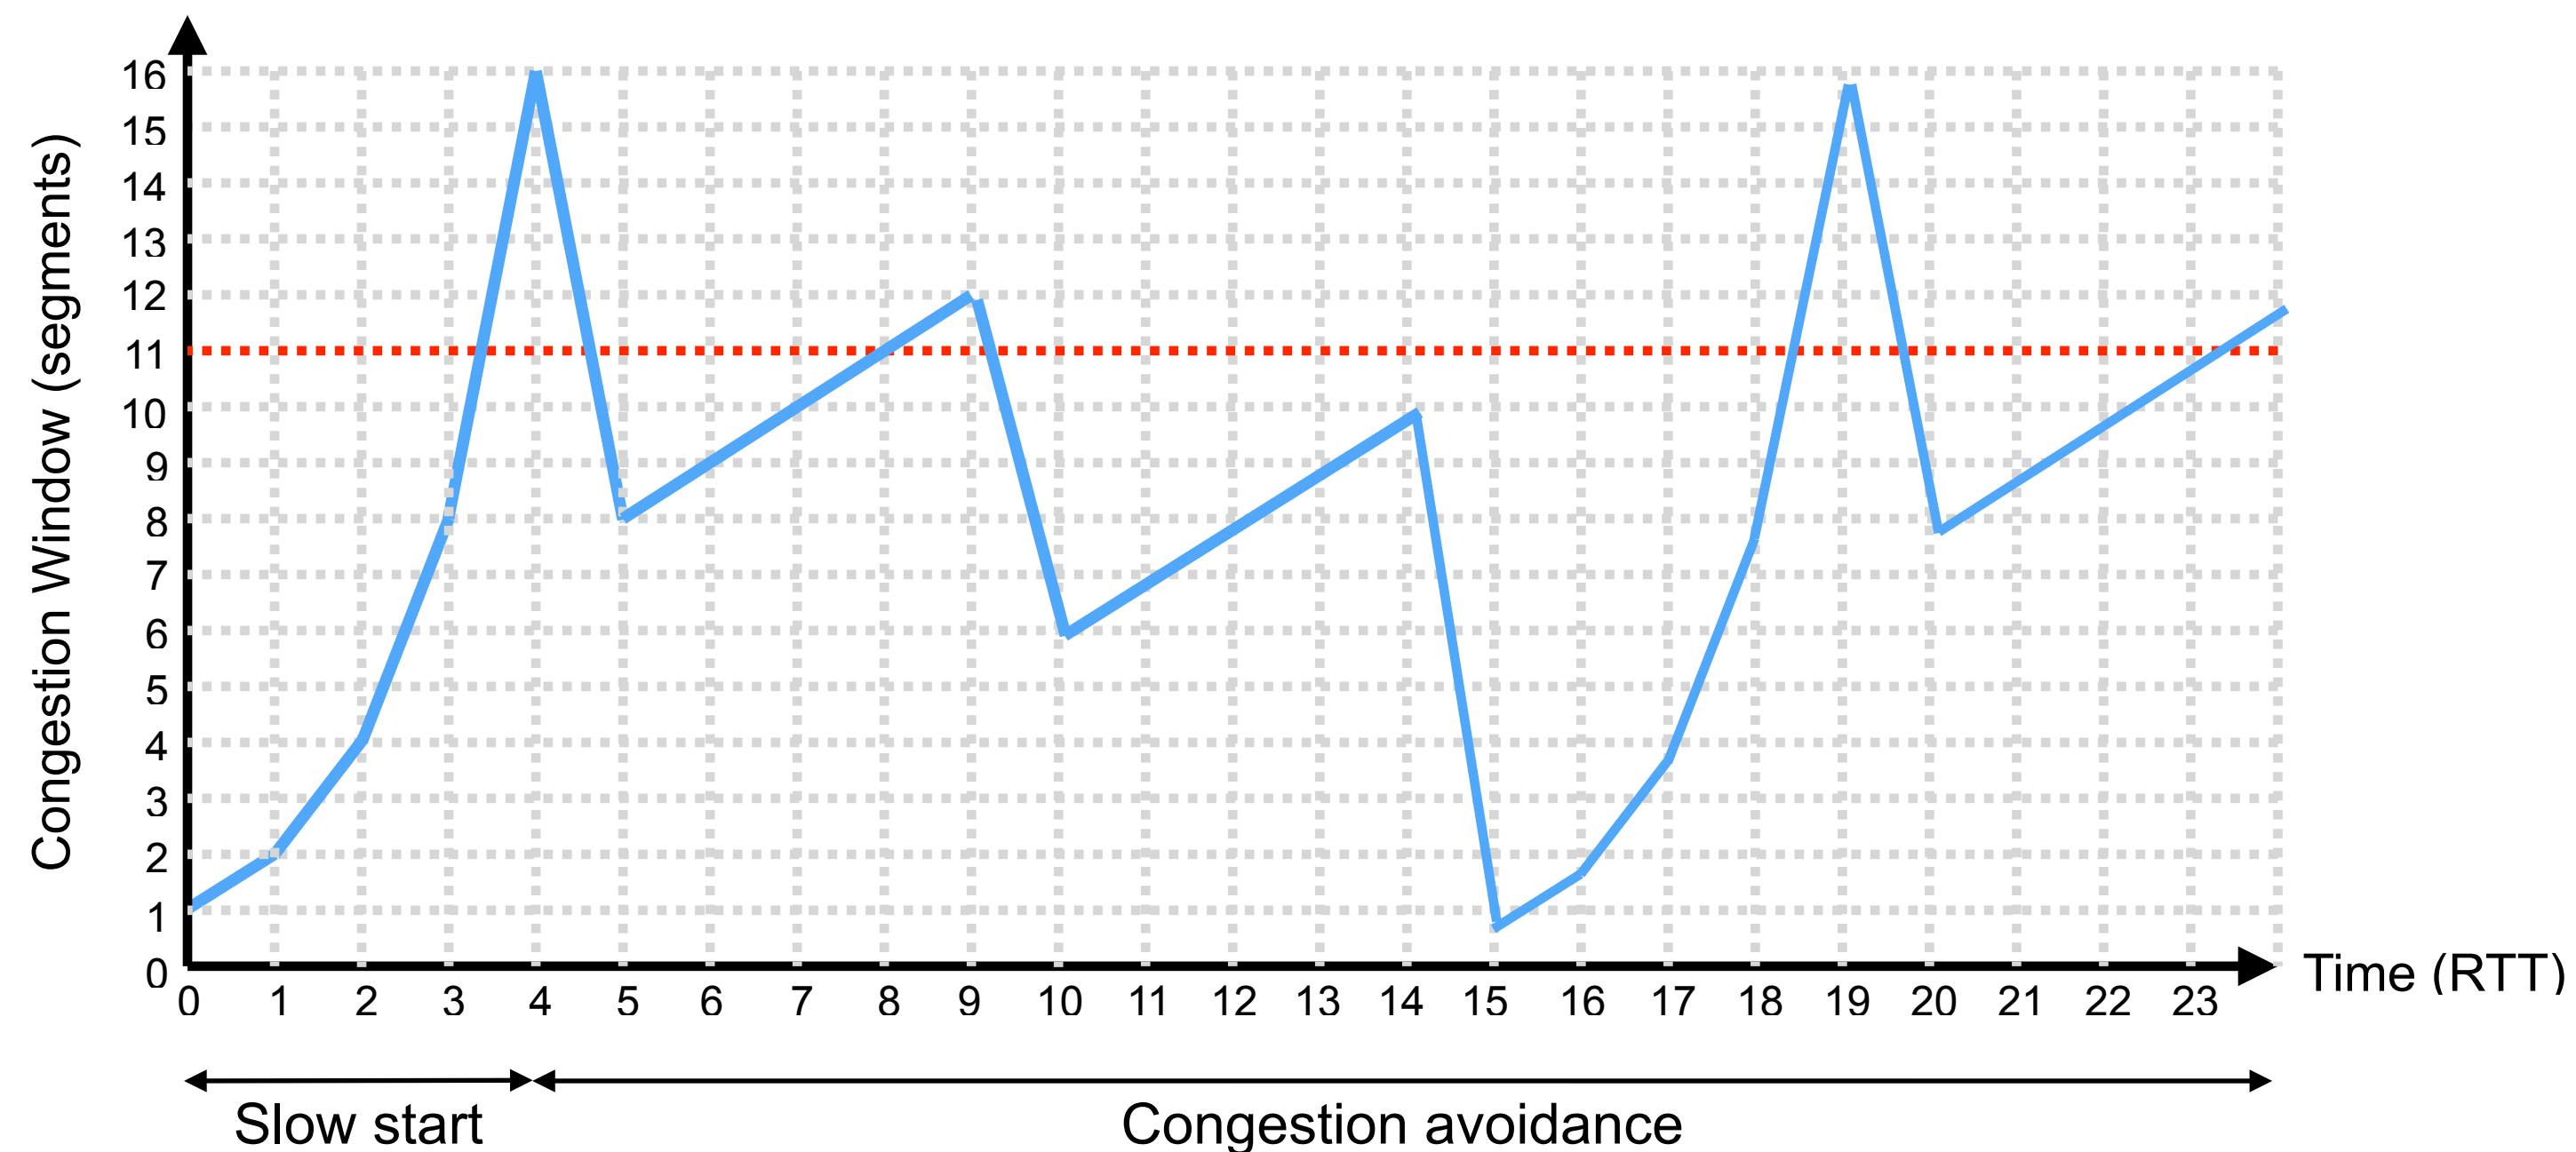
TCP Reno: Congestion Avoidance (3/4)

- If a packet is lost and detected via triple duplicate acknowledgement:
 - Transient congestion, but data still being received
 - Multiplicative decrease in window: $W_i = W_{i-1} \times 0.5$
 - Rapid reduction in window allows congestion to clear quickly, avoids congestion collapse
- Then, return to additive increase until next loss





TCP Reno: Congestion Avoidance (4/4)

- If a packet is lost and detected via timeout:
 - Either receiver or path has failed – reset to W_{init} and re-enter slow start
 - How long is the timeout?
 - $T_{rto} = \max(1 \text{ second, average RTT} + (4 \times \text{RTT variance}))$



Congestion Window Growth, Buffering, Throughput



Bottleneck buffer size = bandwidth \times delay

- Bottleneck queue never empty
- Bottleneck link never becomes idle \rightarrow sending rate varies, but receiver sees continuous flow
- Congestion window follows a “sawtooth” pattern, but received rate is constant at approximately bottleneck bandwidth

Source: G. Appenzeller, “Sizing Router Buffers”, PhD thesis, Stanford University, March 2005.
<http://tiny-tera.stanford.edu/~nickm/papers/guido-thesis.pdf> (Figures 2.1 and 2.2)

TCP Reno: Discussion

- TCP Reno is effective at keeping bottleneck link fully utilised
 - Trades some extra delay to maintain throughput
 - Provided sufficient buffering in the network: $\text{buffer size} = \text{bandwidth} \times \text{delay}$
 - Packets queued in buffer \rightarrow delay
- Limitations:
 - Assumes packet loss is due to congestion; non-congestive loss, e.g., due to wireless interference, impacts throughput
 - Congestion avoidance phase takes a long time to use increased capacity

TCP Reno

- Basic TCP congestion control
- Sliding window algorithms
- Slow start
- Congestion avoidance