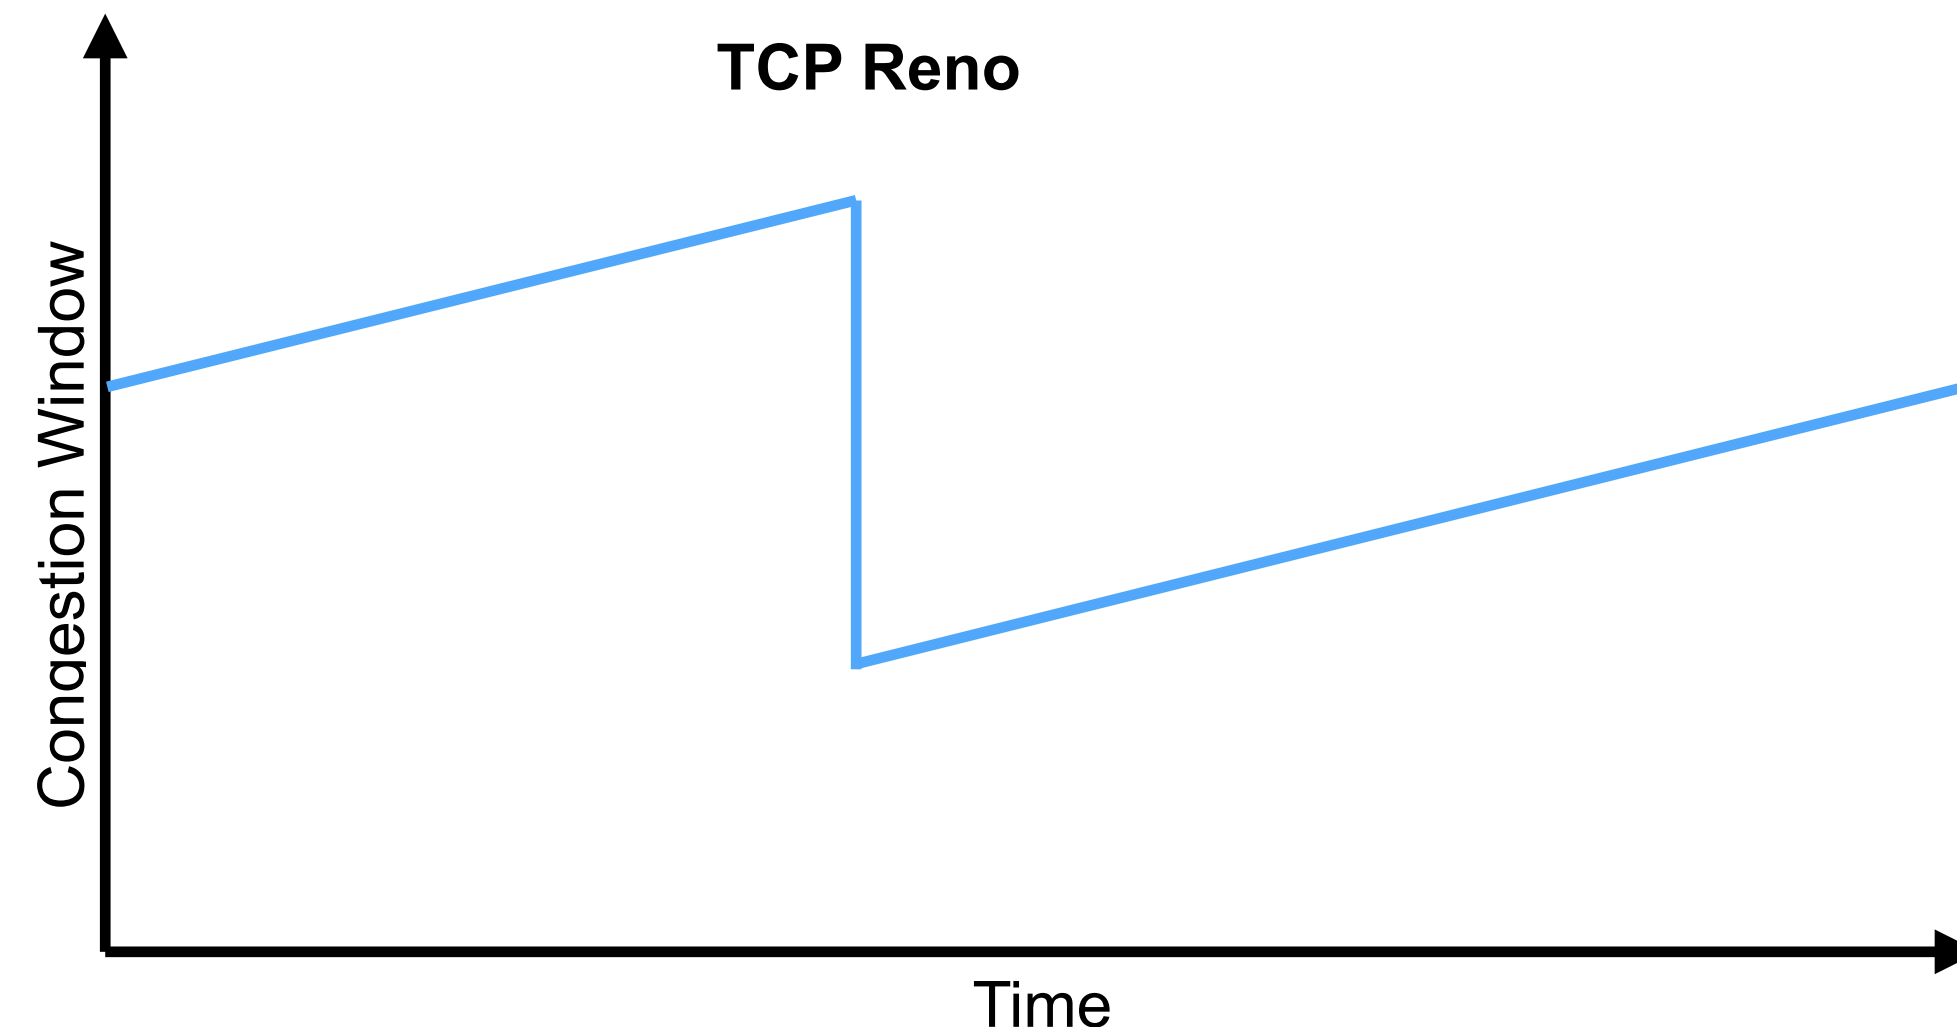


TCP Cubic

- Improving TCP performance on fast, long-distance, networks

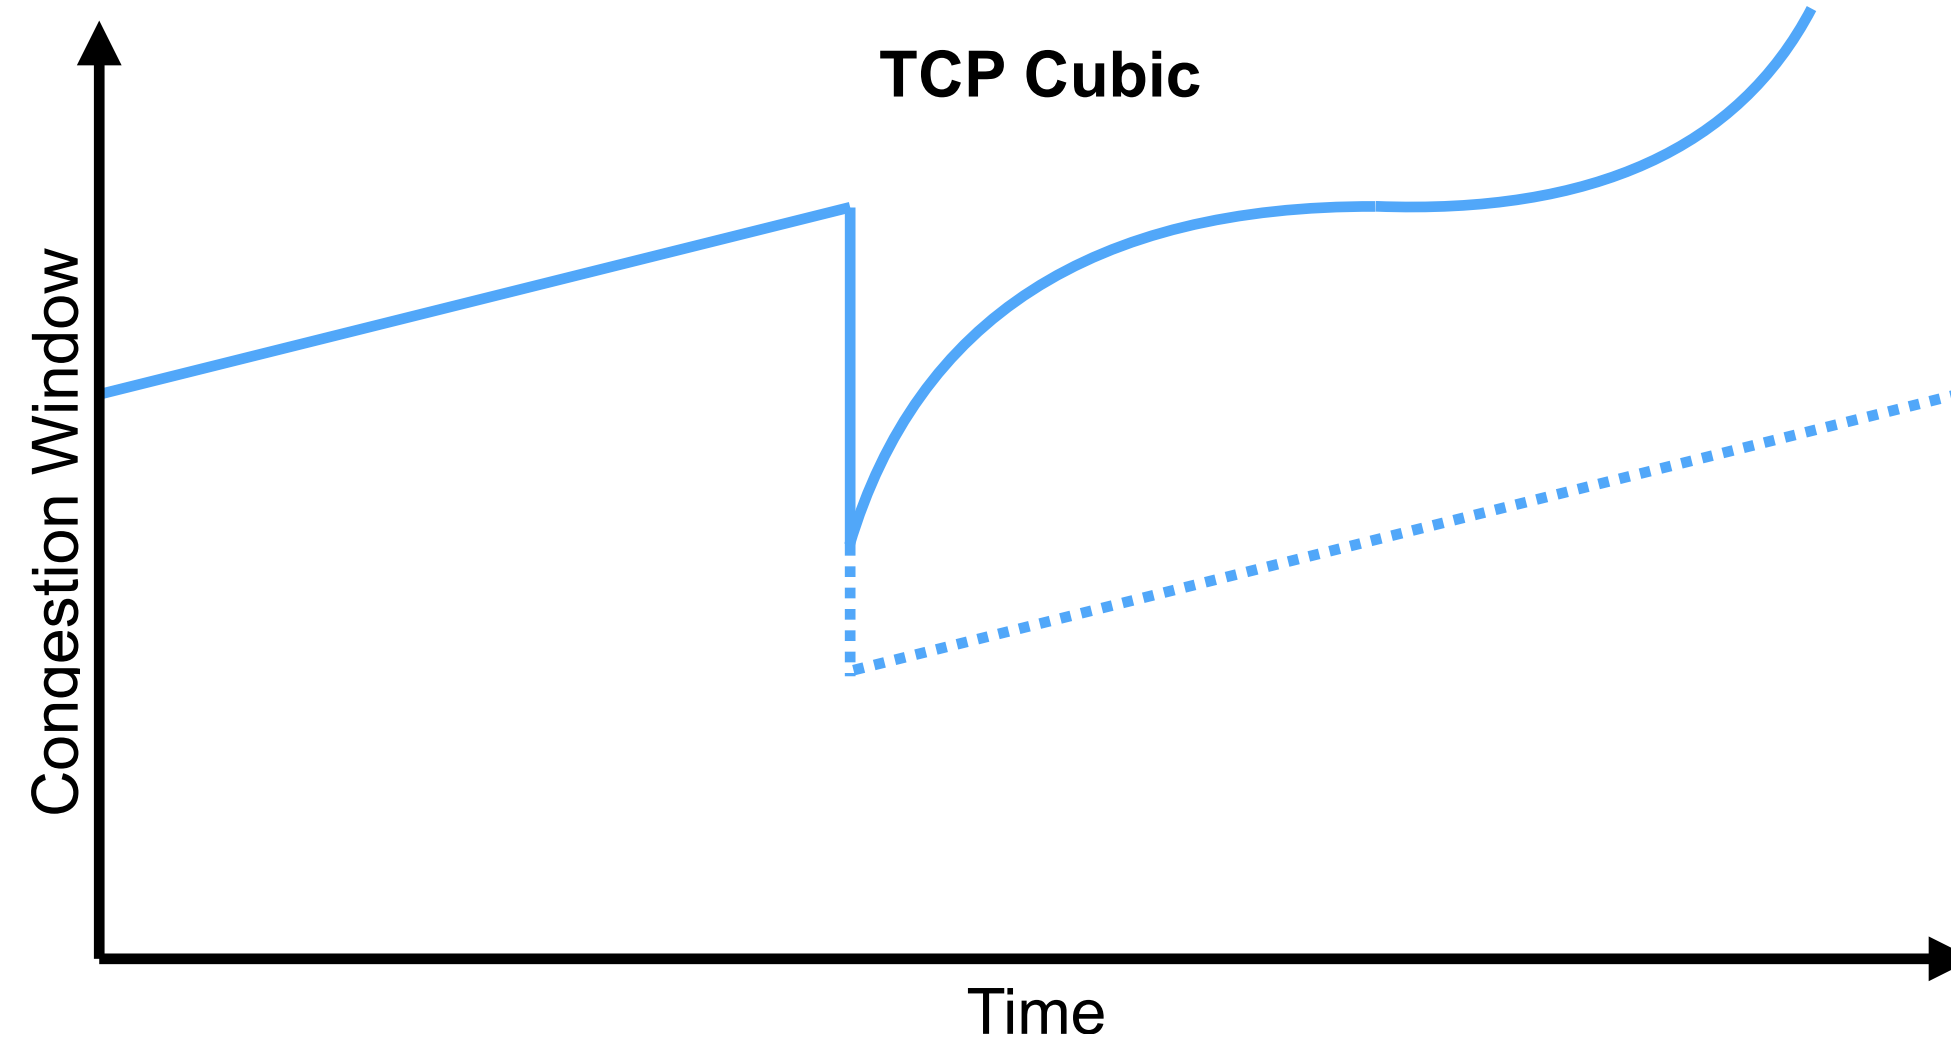
TCP Performance on Fast Long-distance Networks

- TCP Reno can perform poorly on **fast long-distance networks**
 - e.g., multi-gigabit per second inter-continental links
 - Path with 10Gbps bandwidth and 100ms RTT requires congestion window $\sim 100,000$ packets
 - In congestion avoidance, one packet lost and detected via triple duplicate ACK halves the window, then increase by 1 packet per RTT $\rightarrow 50,000$ RTTs to recover sending rate
 - Approximately 1.4 hours with 100ms RTT!



TCP Cubic

- **TCP Cubic** changes the congestion control algorithm
 - During congestion avoidance, increases congestion window faster than TCP Reno on fast long-distance networks
 - Rapidly increases congestion window after packet loss
 - Slows rate of increase as window approaches the largest successfully achieved window



CUBIC: A New TCP-Friendly High-Speed TCP Variant

Sangtae Ha, Injong Rhee
Dept of Computer Science
North Carolina State University
Raleigh, NC 27695
(sha2,rhee)@ncsu.edu

Lisong Xu
Dept of Comp. Sci. and Eng.
University of Nebraska
Lincoln, Nebraska 68588
xu@cse.unl.edu

ABSTRACT
CUBIC is a congestion control protocol for TCP (transmission control protocol) and the current default TCP algorithm in Linux. The protocol modifies the linear window growth function of existing TCP standards to be a cubic function in order to improve the scalability of TCP over fast and long distance networks. It also achieves more equitable bandwidth allocations among flows with different RTTs (round trip times) by making the window growth to be independent of RTT – thus those flows grow their congestion window at the same rate. During steady state, CUBIC increases the window size aggressively when the window is far from the saturation point, and the slowly when it is close to the saturation point. This feature allows CUBIC to be very scalable when the bandwidth and delay product of the network is large, and at the same time, be highly stable and also fair to standard TCP flows. The implementation of CUBIC in Linux has gone through several upgrades. This paper documents its design, implementation, performance and evolution as the default TCP algorithm of Linux.

1. INTRODUCTION
As the Internet evolves to include many very high speed and long distance network paths, the performance of TCP was challenged. These networks are characterized by large bandwidth and delay product (BDP) which represents the total number of packets needed in flight while keeping the bandwidth fully utilized, in other words, the size of the congestion window. In standard TCP like TCP-Reno, TCP-NewReno and TCP-SACK, TCP grows its window one per round trip time (RTT). This makes the data transport speed of TCP used in all major operating systems including Windows and Linux rather sluggish, to say the least, extremely under-utilizing the networks especially if the length of flows is much shorter than the time TCP grows its windows to the full size of the BDP of a path. For instance, if the bandwidth of a network path is 10 Gbps and the RTT is 100 ms, with packets of 1250 bytes, the BDP of the path is around 100,000 packets. For TCP to grow its window from the mid-point of the BDP, say 50,000, it takes about 50,000 RTTs which amounts to 5000 seconds (1.4 hours). If a flow finishes before that time, it severely under-utilizes the path.

To counter this under-utilization problem of TCP, many
¹A short version [27] of this paper was presented at the International Workshop on Protocols for Fast and Long Distance Networks in 2005.
²For brevity, we also denote *Standard TCP* as *TCP*.

64

S. Ha, I. Rhee, and L. Xu, "CUBIC: A New TCP-Friendly High-Speed TCP Variant", ACM SIGOPS Operating System Review, Vol. 54, Num. 5, pages 64-74, July 2008
<https://dx.doi.org/10.1145/1400097.1400105>

I. Rhee, L. Xu, S. Ha, A. Zimmermann, L. Eggert, R. Scheffenegger, "CUBIC for Fast Long-Distance Networks", IETF, February 2018, RFC 8312.
<https://datatracker.ietf.org/doc/rfc8312/>

TCP Cubic Congestion Control

- **On packet loss** during congestion avoidance, TCP Cubic reduces congestion window: $W_i = W_{i-1} \times 0.7$
- TCP Reno uses $W_i = W_{i-1} \times 0.5$
- TCP Cubic is more aggressive
- **After packet loss** during congestion avoidance, TCP Cubic increases the congestion window as:
 $W_{\text{cubic}} = C (t - K)^3 + W_{\text{max}}$
 - W_{max} is the maximum window size reached before the loss
 - t is the time since the packet loss
 - K is the time it will take to increase the window back to W_{max} , assuming no further packet losses
 - $C = 0.4$ is a constant that controls fairness to TCP Reno
- **Many** additional details included to ensure fairness with TCP Reno on slower, shorter-RTT, networks

S. Ha, I. Rhee, and L. Xu, "CUBIC: A New TCP-Friendly High-Speed TCP Variant", ACM SIGOPS Operating System Review, Vol. 54, Num. 5, pages 64-74, July 2008 <https://dx.doi.org/10.1145/1400097.1400105>

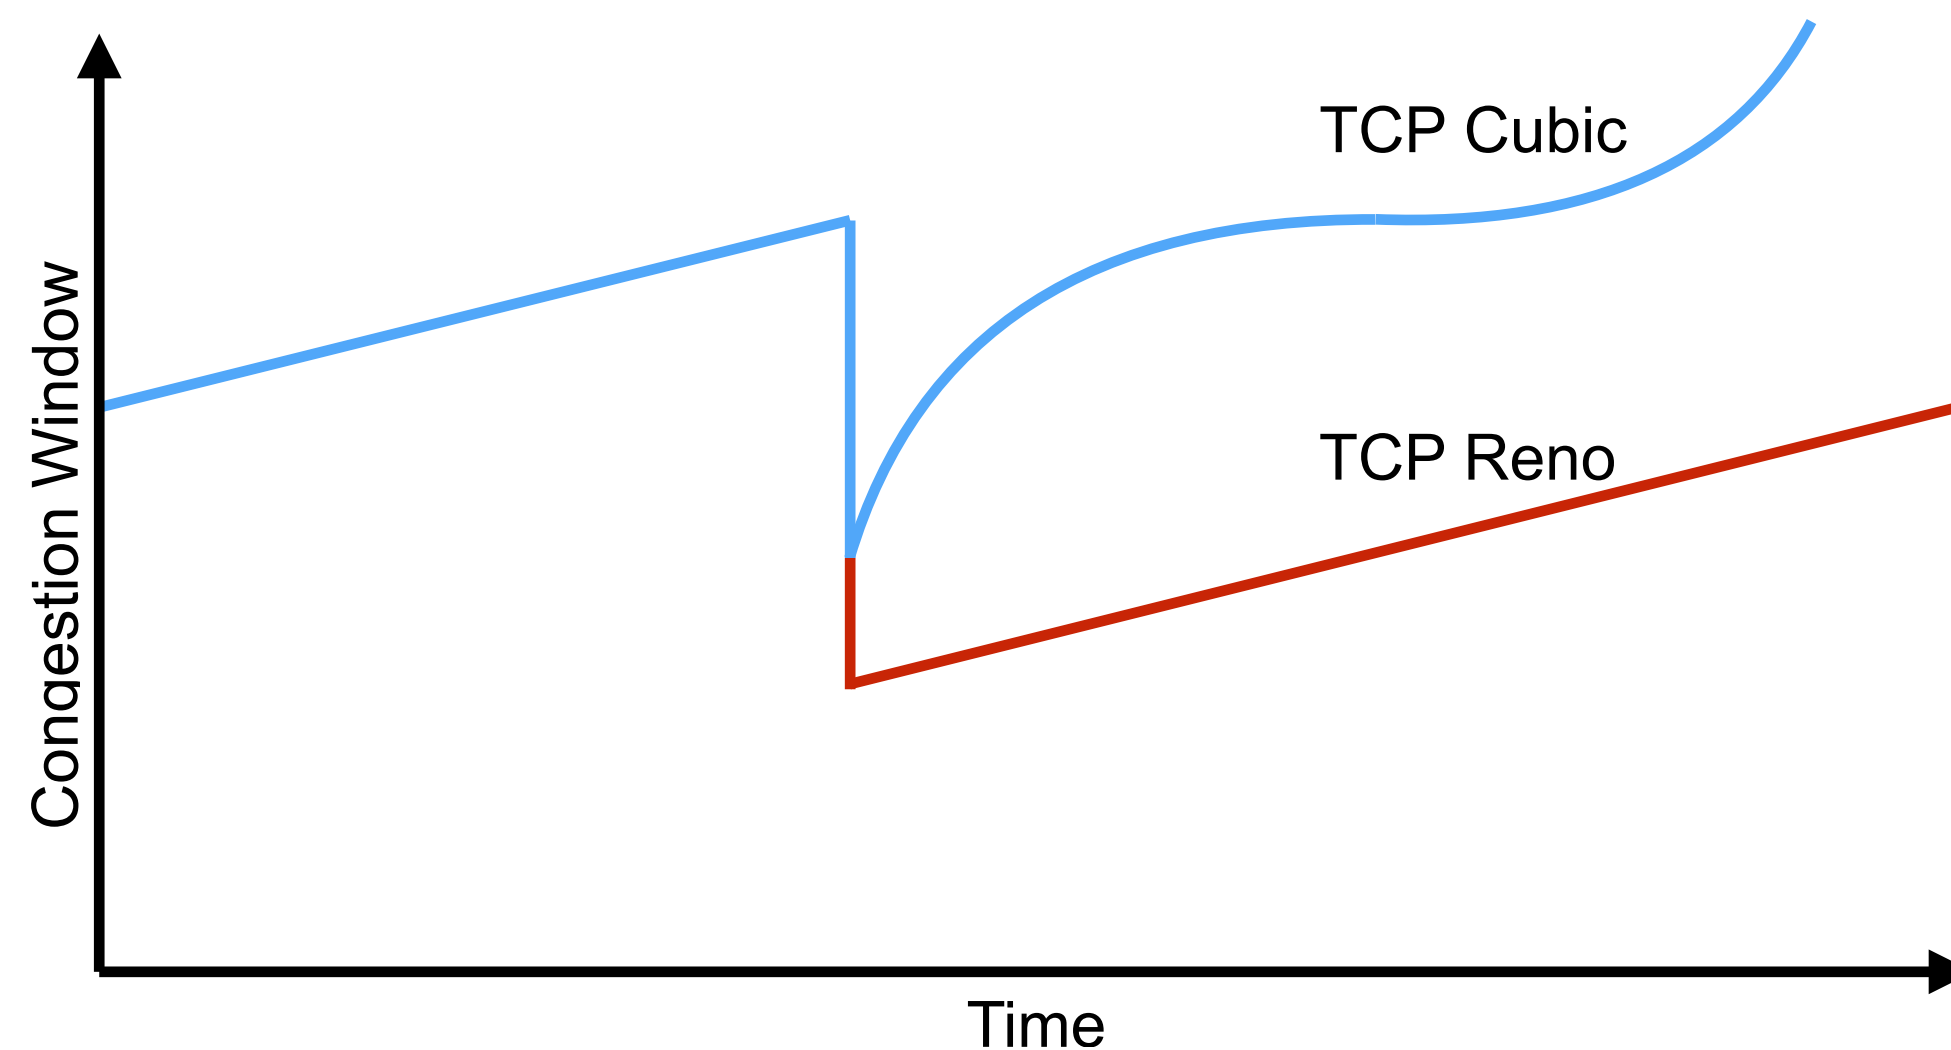
Algorithm 1: Linux CUBIC algorithm (v2.2)

```

Initialization:
  tcp_friendlyness ← 1, β ← 0.2
  fast_convergence ← 1, C ← 0.4
  cubic_reset()
On each ACK:
begin
  if dMin then dMin ← min(dMin, RTT)
  else dMin ← RTT
  if cwnd ≤ ssthresh then cwnd ← cwnd + 1
  else
    cnt ← cubic_update()
    if cwnd_cnt > cnt then
      cwnd ← cwnd + 1, cwnd_cnt ← 0
    else cwnd_cnt ← cwnd_cnt + 1
  end
end
Packet loss:
begin
  epoch_start ← 0
  if cwnd < Wlast_max and fast_convergence then
    Wlast_max ← cwnd *  $\frac{(2-\beta)}{2}$  ..... (3.7)
  else Wlast_max ← cwnd
  ssthresh ← cwnd * (1 - β) ..... (3.6)
end
Timeout:
begin
  cubic_reset()
end
cubic_update(): ..... (3.2)
begin
  ack_cnt ← ack_cnt + 1
  if epoch_start ≤ 0 then
    epoch_start ← tcp_time_stamp
    if cwnd < Wlast_max then
      K ←  $\sqrt[3]{\frac{W_{last\_max} - cwnd}{C}}$ 
      origin_point ← Wlast_max
    else
      K ← 0
      origin_point ← cwnd
    end
    ack_cnt ← 1
    Wtcp ← cwnd
    t ← tcp_time_stamp + dMin - epoch_start
    target ← origin_point + C(t - K)3
    if target > cwnd then cnt ←  $\frac{cwnd}{target - cwnd}$  .. (3.4,3.5)
    else cnt ← 100 * cwnd
    if tcp_friendlyness then cubic_tcp_friendlyness()
  end
  cubic_tcp_friendlyness(): ..... (3.3)
begin
  Wtcp ← Wtcp +  $\frac{3\beta}{2-\beta} * \frac{ack\_cnt}{cwnd}$ 
  ack_cnt ← 0
  if Wtcp > cwnd then
    max_cnt ←  $\frac{cwnd}{W_{tcp} - cwnd}$ 
    if cnt > max_cnt then cnt ← max_cnt
  end
end
cubic_reset():
begin
  Wlast_max ← 0, epoch_start ← 0, origin_point ← 0
  dMin ← 0, Wtcp ← 0, K ← 0, ack_cnt ← 0
end
  
```

TCP Cubic vs Reno

- TCP Cubic is default in most modern operating systems
 - Much more complex than TCP Reno – core response is relatively straight-forward
 - Much complexity ensures fairness with TCP Reno in its typical operating regime; improves performance for networks with longer RTT and higher bandwidth
- Both algorithms use packet loss as congestion signal and eventually fill router buffers
 - Trade latency for throughput



TCP Cubic

- Improving TCP performance on fast, long-distance, networks