

TCP-CUBIC Overview

TCP-CUBIC is an improvement over TCP-BIC in the way that while retaining most of strengths of BIC (especially, its stability and scalability under high BDP networks), it simplifies the window control function, enhances its TCP friendliness and achieves better RTT fairness. CUBIC also performs better under small BDP networks while BIC can still be too aggressive for TCP under short RTT or small BDP networks.

The window growth function of CUBIC is a function of elapsed time since the last packet loss event:

$$W(t) = C(t - K)^3 + W_{max}$$
$$K = \sqrt[3]{\frac{W_{max}\beta}{C}}$$

C is a constant scaling factor with default value of 0.4, β is the multiplicative decrease factor after a packet loss event, its default value is 0.2, and W_{max} is the window size just before the last window reduction. Here K represents the time period needed to increase W to W_{max} .

CUBIC Algorithm

Algorithm 1 TCP CUBIC algorithm for NDN

```
1: procedure INIT() ▷ CUBIC initialization
2:   tcp_friendliness  $\leftarrow$  true,  $\beta \leftarrow 0.2$ 
3:   fast_convergence  $\leftarrow$  true,  $C \leftarrow 0.4$ 
4:   CubicReset()
5:
6: procedure ONDATA() ▷ Data packets act as selective acknowledgement of Interest packets
7:   if min_rtt then min_rtt  $\leftarrow$  min(min_rtt, RTT)
8:   else min_rtt  $\leftarrow$  RTT
9:   if cwnd  $\leq$  ssthresh then
10:    cwnd  $\leftarrow$  cwnd + 1 ▷ slow start
11:   else
12:    CubicUpdate() ▷ congestion avoidance
13:
14: procedure ONPACKETLOSS() ▷ Indicated either by timeouts or NACKs
15:   epoch_start  $\leftarrow$  0
16:   if cwnd  $<$  Wlast_max & fast_convergence then
17:    Wlast_max  $\leftarrow$  cwnd *  $\frac{2-\beta}{2}$  ▷ release more bandwidth for new flows to catch up
18:   else
19:    Wlast_max  $\leftarrow$  cwnd
20:   cwnd  $\leftarrow$  cwnd * (1 -  $\beta$ )
21:   ssthresh  $\leftarrow$  cwnd
22:
23: procedure CUBICUPDATE() ▷ Update congestion window
24:   if epoch_start  $\leq$  0 then
25:    epoch_start  $\leftarrow$  current_time
26:    if cwnd  $<$  Wlast_max then
27:      $K \leftarrow \sqrt[3]{\frac{W_{last\_max} - cwnd}{C}}$  ▷ k is in second
28:     origin_point  $\leftarrow$  Wlast_max
29:    else
30:      $K \leftarrow 0$ 
31:     origin_point  $\leftarrow$  cwnd
32:    Wtcp  $\leftarrow$  cwnd
33:     $t \leftarrow$  current_time + min_rtt - epoch_start ▷ t is in second
34:    target  $\leftarrow$  origin_point +  $C(t - K)^3$  ▷ calculate  $W(t + rtt)$ 
35:    if target  $>$  cwnd then
36:     cwnd_update  $\leftarrow$  cwnd +  $\frac{target - cwnd}{cwnd}$ 
37:    else
38:     cwnd_update  $\leftarrow$  cwnd +  $\frac{0.01}{cwnd}$  ▷ only a small increment
39:    if tcp_friendliness then ▷ make sure window grows at least at the speed of TCP
40:      $W_{tcp} \leftarrow W_{tcp} + \frac{3\beta}{2-\beta} * \frac{t}{RTT}$ 
41:     if  $W_{tcp} > cwnd$  &  $W_{tcp} > target$  then
42:      cwnd_update  $\leftarrow$  cwnd +  $\frac{W_{tcp} - cwnd}{cwnd}$ 
43:    cwnd  $\leftarrow$  cwnd_update ▷ update window size
44:
45: procedure CUBICRESET() ▷ Reset state variables
46:   Wlast_max  $\leftarrow$  0, epoch_start  $\leftarrow$  0, origin_point  $\leftarrow$  0
47:   min_rtt  $\leftarrow$  0, Wtcp  $\leftarrow$  0,  $K \leftarrow 0$ , ssthresh  $\leftarrow$  MAX_INT
```

Implementation Challenges

The implementation we can refer to is the one implemented in Linux kernel (http://lxr.linux.no/linux/net/ipv4/tcp_cubic.c). It's quite hard to get because it uses a bunch of scaling factors and performs complex unit conversions. I think it's because Linux kernel doesn't allow floating point operations. Since we are going to implement CUBIC in user mode, it's not a problem. But the concern is that floating point operations may introduce numerical stability. At the bottom line, we can still mimic the kernel's implementation by using scaling factors.

Test Scenarios

Topology	Bottleneck Link Bandwidth	Flows	Purposes
Linear topology	5Mbps	single CUBIC flow	how CUBIC performs under low BDP network
Linear topology	100Mbps	single CUBIC flow	how CUBIC performs under high BDP network
Dumbbell topology	100Mbps	two CUBIC flows	see how fast two flow converges
Dumbbell topology	100Mbps	one CUBIC flow and one TCP-Reno flow	test intra-protocol fairness

References

Original CUBIC paper: http://netsrv.csc.ncsu.edu/export/cubic_a_new_tcp_2008.pdf
CUBIC implementation in ns-3: http://web.cs.wpi.edu/~rek/Research/Papers/WNS3_14.pdf
CUBIC implementation in Linux kernel: http://lxr.linux.no/linux/net/ipv4/tcp_cubic.c