XCP

Thursday, February 8, 2018 12:45 PM

XCP Motivation

1.) Large BDP environments result in poor TCP performance.

2.) Why? (Recall from Handley's paper)

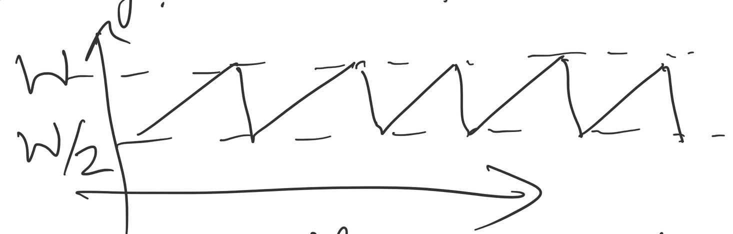

Steady state throughput of TCP

$= \frac{\text{Avg. window size}}{RTT}$ (Ignore queue delay in RTT)

$= \frac{3W}{4RTT}$

Window goes from $\frac{W}{2}$ to W in time Δt.

What's Δt?

Additive increase implies window goes from $\frac{W}{2}$ to W in time $\frac{W}{2}$ RTT's

So the first RTT we send $\frac{W}{2}$ packets

second RTT $\frac{W}{2} + 1$

third RTT $\frac{W}{2} + 2$

⋮

$\frac{W}{2}$ RTT $\frac{W}{2} + \frac{W}{2}$

Total number of packets

$\approx \frac{W}{2} * \frac{W}{2} + \frac{W}{2} * \frac{W}{4}$

$\approx \frac{W^2}{4} + \frac{W^2}{8}$

$\approx \frac{3W^2}{8}$

Loss probability $= \frac{1}{\frac{3W^2}{8}}$

$p = \frac{8}{3W^2}$

$W = \sqrt{\frac{8}{3p}}$

Plugging back into throughput

$Tpt = \frac{3W}{4RTT}$

$= \frac{3}{4}\sqrt{\frac{8}{3p}} \times \frac{1}{RTT}$

$$\boxed{Tpt = \sqrt{\frac{3}{2p}} \times \frac{1}{RTT}}$$

What's the point of this?

1.) Tpt ↓ as RTT ↑ (So large RTT flows suffer)

2.) W needs to be close to B*RTT support full utilization.

$p = \frac{8}{3W^2}$

Hence p needs to be $\propto \frac{1}{B^2 RTT^2}$ for full utilization.

1.) But p typically goes up with RTT (more hops, more chance to drop packets)

2.) & remains at most constant with B (larger B needs larger buffer size to sustain larger incoming packet rate)

$\boxed{\text{Net effect: TCP does badly when BDP} \uparrow}$