BBRv2 Update:

QUIC Tweaks and Internet Deployment

TCP BBR: Neal Cardwell, Yuchung Cheng, Kevin Yang

Soheil Hassas Yeganeh, Priyaranjan Jha, Yousuk Seung, Luke Hsiao, Matt Mathis

Van Jacobson

QUIC BBR: Ian Swett, Bin Wu, Victor Vasiliev

https://groups.google.com/d/forum/bbr-dev

Google

IETF 112: Online; Nov 8, 2021

Outline

The following tweaks have gotten QUIC's BBRv2 close to BBRv1 Video QoE and on par for Search Latency

Core issues:

- Setting inflight_hi
- Early PROBE_UP exit
- Excessive time in PROBE_RTT

Setting inflight_hi

Exiting STARTUP due to loss:

```
        Inflight_hi = BDP();
```

Setting inflight_hi to BDP => likely to be CWND limited before max bandwidth

Once inflight_hi is low, it can be difficult or impossible to grow (see later)

Setting inflight_hi

Exiting STARTUP due to loss:

Inflight_hi = max(BDP(), **max_delivered_in_round**)

Bytes delivered in a round indicates the pipe is at least that large

Less bandwidth crash with aggregation

Early PROBE_UP exit

PROBE_UP can exit early due to ‘queuing’

Exit if:

In PROBE_UP for at least min_rtt AND bytes_in_flight >= 1.25 * BDP() + 2*MSS

If you’re not in PROBE_UP, you can’t increase inflight_hi

If inflight_hi doesn’t increase, you may never achieve original max bandwidth

Early PROBE_UP exit

Exit if:

In PROBE_UP for at least **1 round** AND
bytes_in_flight >= 1.25 * BDP() + 2*MSS + **extra_acked**

Avoids immediately exiting PROBE_UP in the presence of aggregation

Early PROBE_UP exit

Newer Idea

Instead of always adding extra_acked, what about checking for a persistent queue?

Exit if:

In PROBE_UP for at least **1 round** AND
min_bytes_in_flight_in_round > 1.25 * BDP() + 2*MSS

Allows skipping the app-limited check

Can also be used for STARTUP exit (code)

*Insignificant application data so far

Excessive Time in PROBE_RTT

Problem: PROBE_RTT limits CWND to ½ BDP, so <½ the bandwidth

Observation: Flows that go idle in PROBE_RTT come out in PROBE_RTT for one round trip

Solution: Exit PROBE_RTT upon exiting idle if the elapsed time is large enough

Avoids an extra round trip in PROBE_RTT

Turns out TCP independently landed this fix!