


Running Code Tests of AQM




Jim Gettys & Dave Taht
IETF 88

Time Domain Behavior is Critical

- Studying steady state behavior is not useful at the edge of the network, where load changes rapidly. Critical interesting questions include:
 - How fast can applications adjust to changes in load?
 - What does RTT's on different paths do? (TCP itself is not fair, and the world is big)
 - Total latency (latency + Jitter) imposed on real time applications caused by competing saturating load: i.e. how well will RTCWeb work on a loaded edge link? How well will VOIP work? How often will I be fragged? Will I make the next stock trade before the other guy?
 - How fast can your load adjust to rapid changes in throughput: wireless can change 10x in 100ms? (future)

"Netalyzr: Illuminating Edge Network Neutrality, Security, and Performance"


C. Kreibich, N. Weaver, B. Nechaev, and V. Paxson



Green diagonal line == .5 second latency

Inferred Buffer Capacity

Uplink



black diagonal line == 4 second latency

Inferred Buffer Capacity

Downlink

Arrow direction is increasing latency

Note: telephony standards for latency are only 150ms!!!

This data is a *lower* bound on the severity of the broadband bufferbloat problem/
This data mixes wireless and wired traffic, so may be contaminated with WiFi bufferbloat


Time Domain Behavior Test

- Linux Netem network emulator discipline: new fixed version on a dedicated box
- 25Mbps down, 25Mbps up
- 4MB of bloat
- Each RRUL test (rtt_fairbe) is 2 streams at 240ms RTT, 2 streams local RTT
- Add a RRUL test @ T+20 and T+40, so at peak, there are 12 bulk TCP streams running flat out up and down, 12 measurement streams
 - BitTorrent typically runs 10-12 less intense flows
- Compared no AQM, PIE v3, fq_codel


No AQM, the usual bloated disaster

RTT Fair Realtime Response Under Load
Download, upload, ping (scaled versions)

download=25000-upload=25000-noecn-bfifo_limit_4096k-2-target=5ms-delay=240ms-quantum_1514-120



Lunar distance




Local/remote: dac/172.21.39.2 - Time: 2013-10-31 01:17:02.139238 - Length/step: 120s/0.20s

Cannot even use the capacity of the link! Wild variation in stream throughput!

PIE response under load


RTT Fair Realtime Response Under Load
Download, upload, ping (scaled versions)



Under load, your VOIP latency + jitter is > 60ms. You will be fragged right and left or lose your shirt in the stock market.

fq_codel response under load

RTT Fair Realtime Response Under Load
Download, upload, ping (scaled versions)




Local/remote: dac/172.21.39.2 - Time: 2013-10-30 23:58:30.783534 - Length/step: 120s/0.20s

The latency under load for sparse traffic is almost immeasurable; zero jitter for the measurement stream @ 25Mbps; so your voip, WebRTC, gaming, DNS, TCP opens, will always fly through even under high load. You win your next Black Ops 2 game, and make a fortune on the market!

3 hop WiFi mesh network under daily load up to hundreds of users

2001:470:1f04:b77::2




IPv6 tunnel from the internet through 3 hops of WiFi through CeroWrt, over 6 weeks

Linux Debloat Highlights

- Codel/fq_codel entered Linux 3.5, 21 July, 2012
- Linux 3.10(?)
 - Netem (network emulator) improvements; should still be on a separate box
- Linux 3.12 (shipped day before yesterday)
 - new “paced fq” scheduler for server TCP
 - TSO sizing: <http://lwn.net/Articles/564978/>
 - Sysctl switch (one line to switch qdiscs on all interfaces)
 - PIE submitted to Linux (probably Linux 3.13?): but no running code experience; 2 major algorithmic versions since May, 20ms target (wtf?)
 - ATM over DSL compensation fixed at last after 18 months being completely broken!

fq_codel + HTB with DSL ATM Compensation

Realtime Response Under Load
Ping CDF plot



Local/remote: alrua-desktop/demo.tohojo.dk - Time: 2013-10-23 13:35:18.681875 - Length/step: 60s/0.20s

Unloaded RTT to this site is 40ms

fq_codel running code field testing



- fq_codel: no tuning needed, first do no harm
 - 5ms target + flow queuing + sparse flow preference
- OpenWrt
 - fq_codel became the default queue discipline last May on all interfaces (users get it **without** asking)
 - People are happy! No known complaints....
 - Next step: fixing WiFi
- Present in DDWrt, CeroWrt, IPFire, Sophos , all current Linux distributions: e.g. Ubuntu, Debian, Fedora, etc...
- A component of StreamBoost(tm)

CeroWrt 3.10.17-6 released

- Think: advanced build of OpenWrt for aqm and other research (e.g. mesh, naming, ...)
 - fq_codel on all wireless and network interfaces
 - Pie available for testing
 - Multiple variants of codel and fq_codel under test
 - Various shapers and tons of test tools
 - Quagga w. OSPF, BGP, ISIS, Babel
- <http://www.bufferbloat.net/projects/cerowrt>
- \$100 home router... Come play and help!

Cable News

- PIE variant mandated baseline in DOCSIS 3.1 cable modems: see the spec for details
- Allows for other AQM and flow queuing
- AQM mandated, but not specified for CMTS's
- Arris has published in SCTE proceedings enhancements to SFQ and LRED

Divisions of Attitudes In this group

- AQM on the core, forget the queues!
- AQM & Packet scheduling on the edge: order of delivery of each packet really matters!

Remember, 1 **packet@1Mbps** == 13ms

- See: <http://gettys.wordpress.com/.....>
- What's right for wireless?
 - Problems w. txops, packet aggregation & diffserv
 - Do we have people here who care and understand about wireless?

Help Wanted

- PIE/fq_codel etc. need integration into NS2 and NS3 code bases to be easy to use by everyone
- CeroWrt needs developers & maintainers
- RRUL tests
- Other Tests like tcptrace need love
- bufferbloat.net site love requested
- One way measurement tests
- WebRTC tests
- How to dispel myths based on >20 year old data?
- New data based on **current** code