Cake

sch_cake

Comprehensive Queue Management Made Easy

Jonathan Morton

Not a Mesh Protocol

- It's a Queue Discipline (`qdisc`) - or AQM.
- Use it with any protocol you like.
  - ...even standard static routing.
- Cake doesn't care.

Not Specialised

- Designed with wired edge nodes in mind.
  - ADSL
  - Cable
  - Fibre
- Not specialised for WiFi or LANs...
  - ...but works anyway!

Baseline

- Fq_codel is in Linux kernel mainline.
  - Codel "controlled delay" AQM
  - DRR++ flow isolation
- HTB is the de-facto standard shaper.
  - Some people use HFSC instead.
- Combination requires expert setup.

AQM

- Unmanaged queue induces large delays...
  - Routinely several seconds...
    - Occasionally several minutes!
- Keep queue lengths short.
- Minimise induced delay under load.
  - Routinely around 10ms within flow.

Life Without AQM

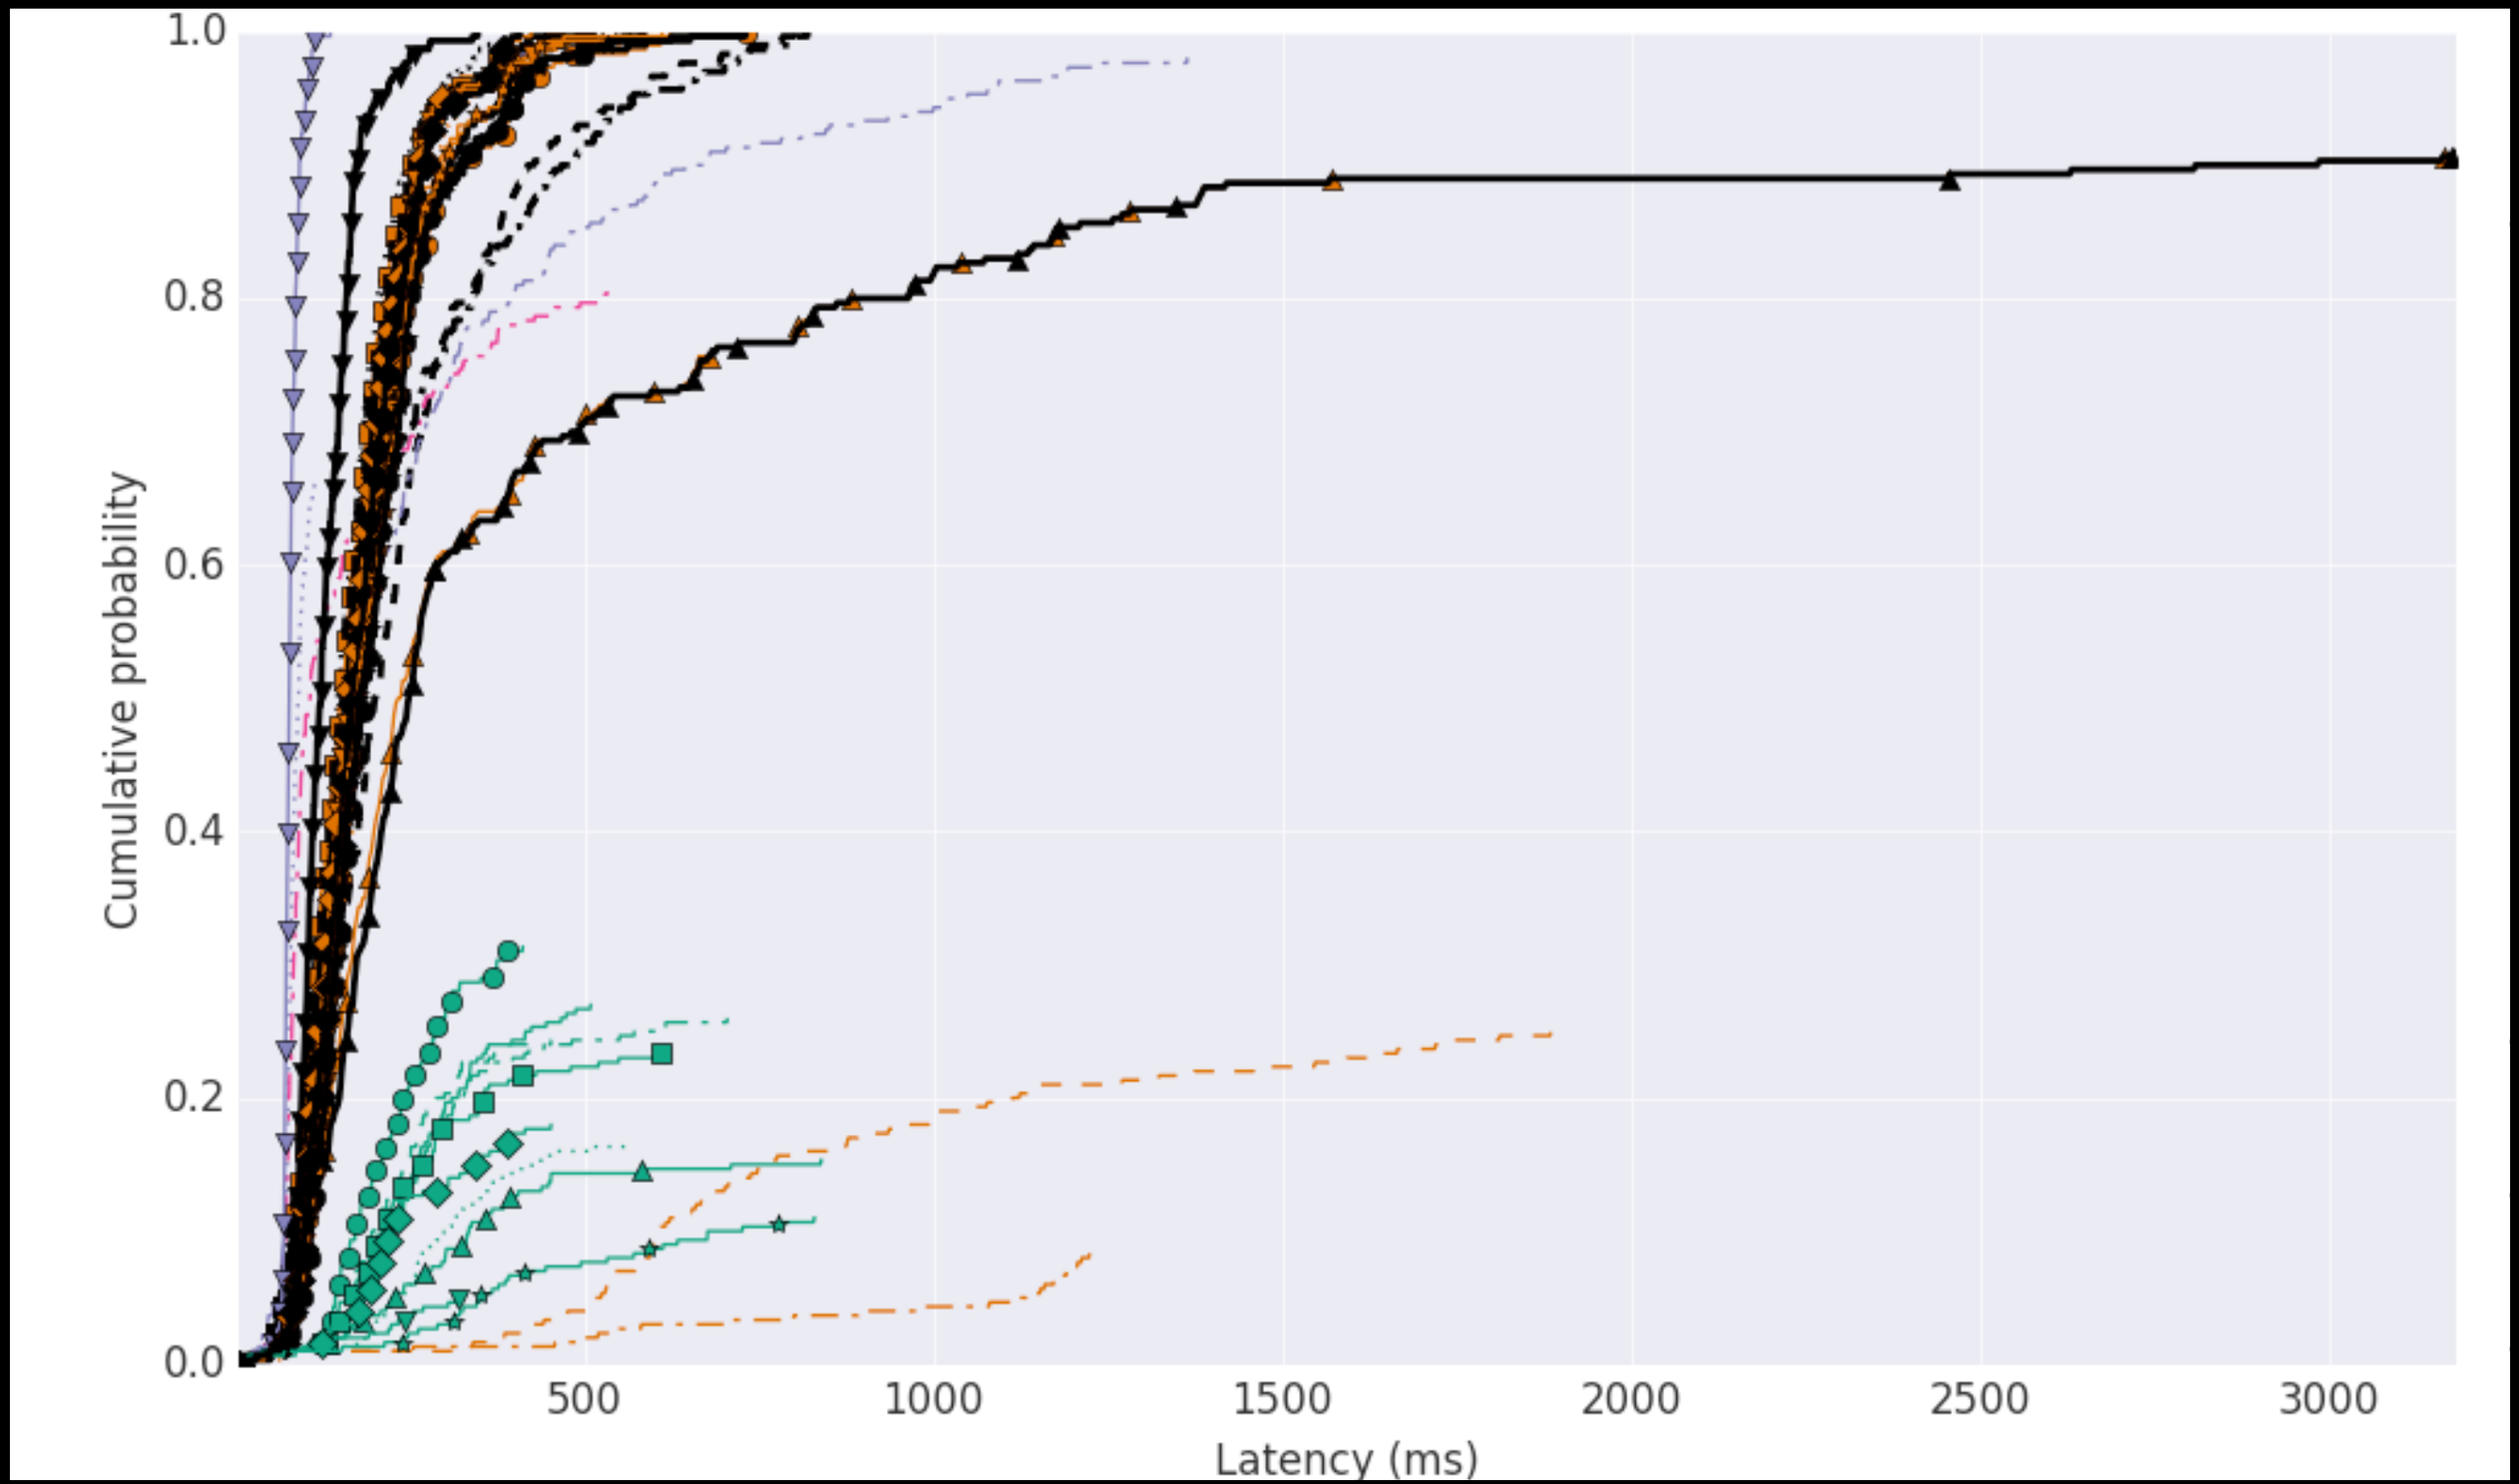

Life Without AQM

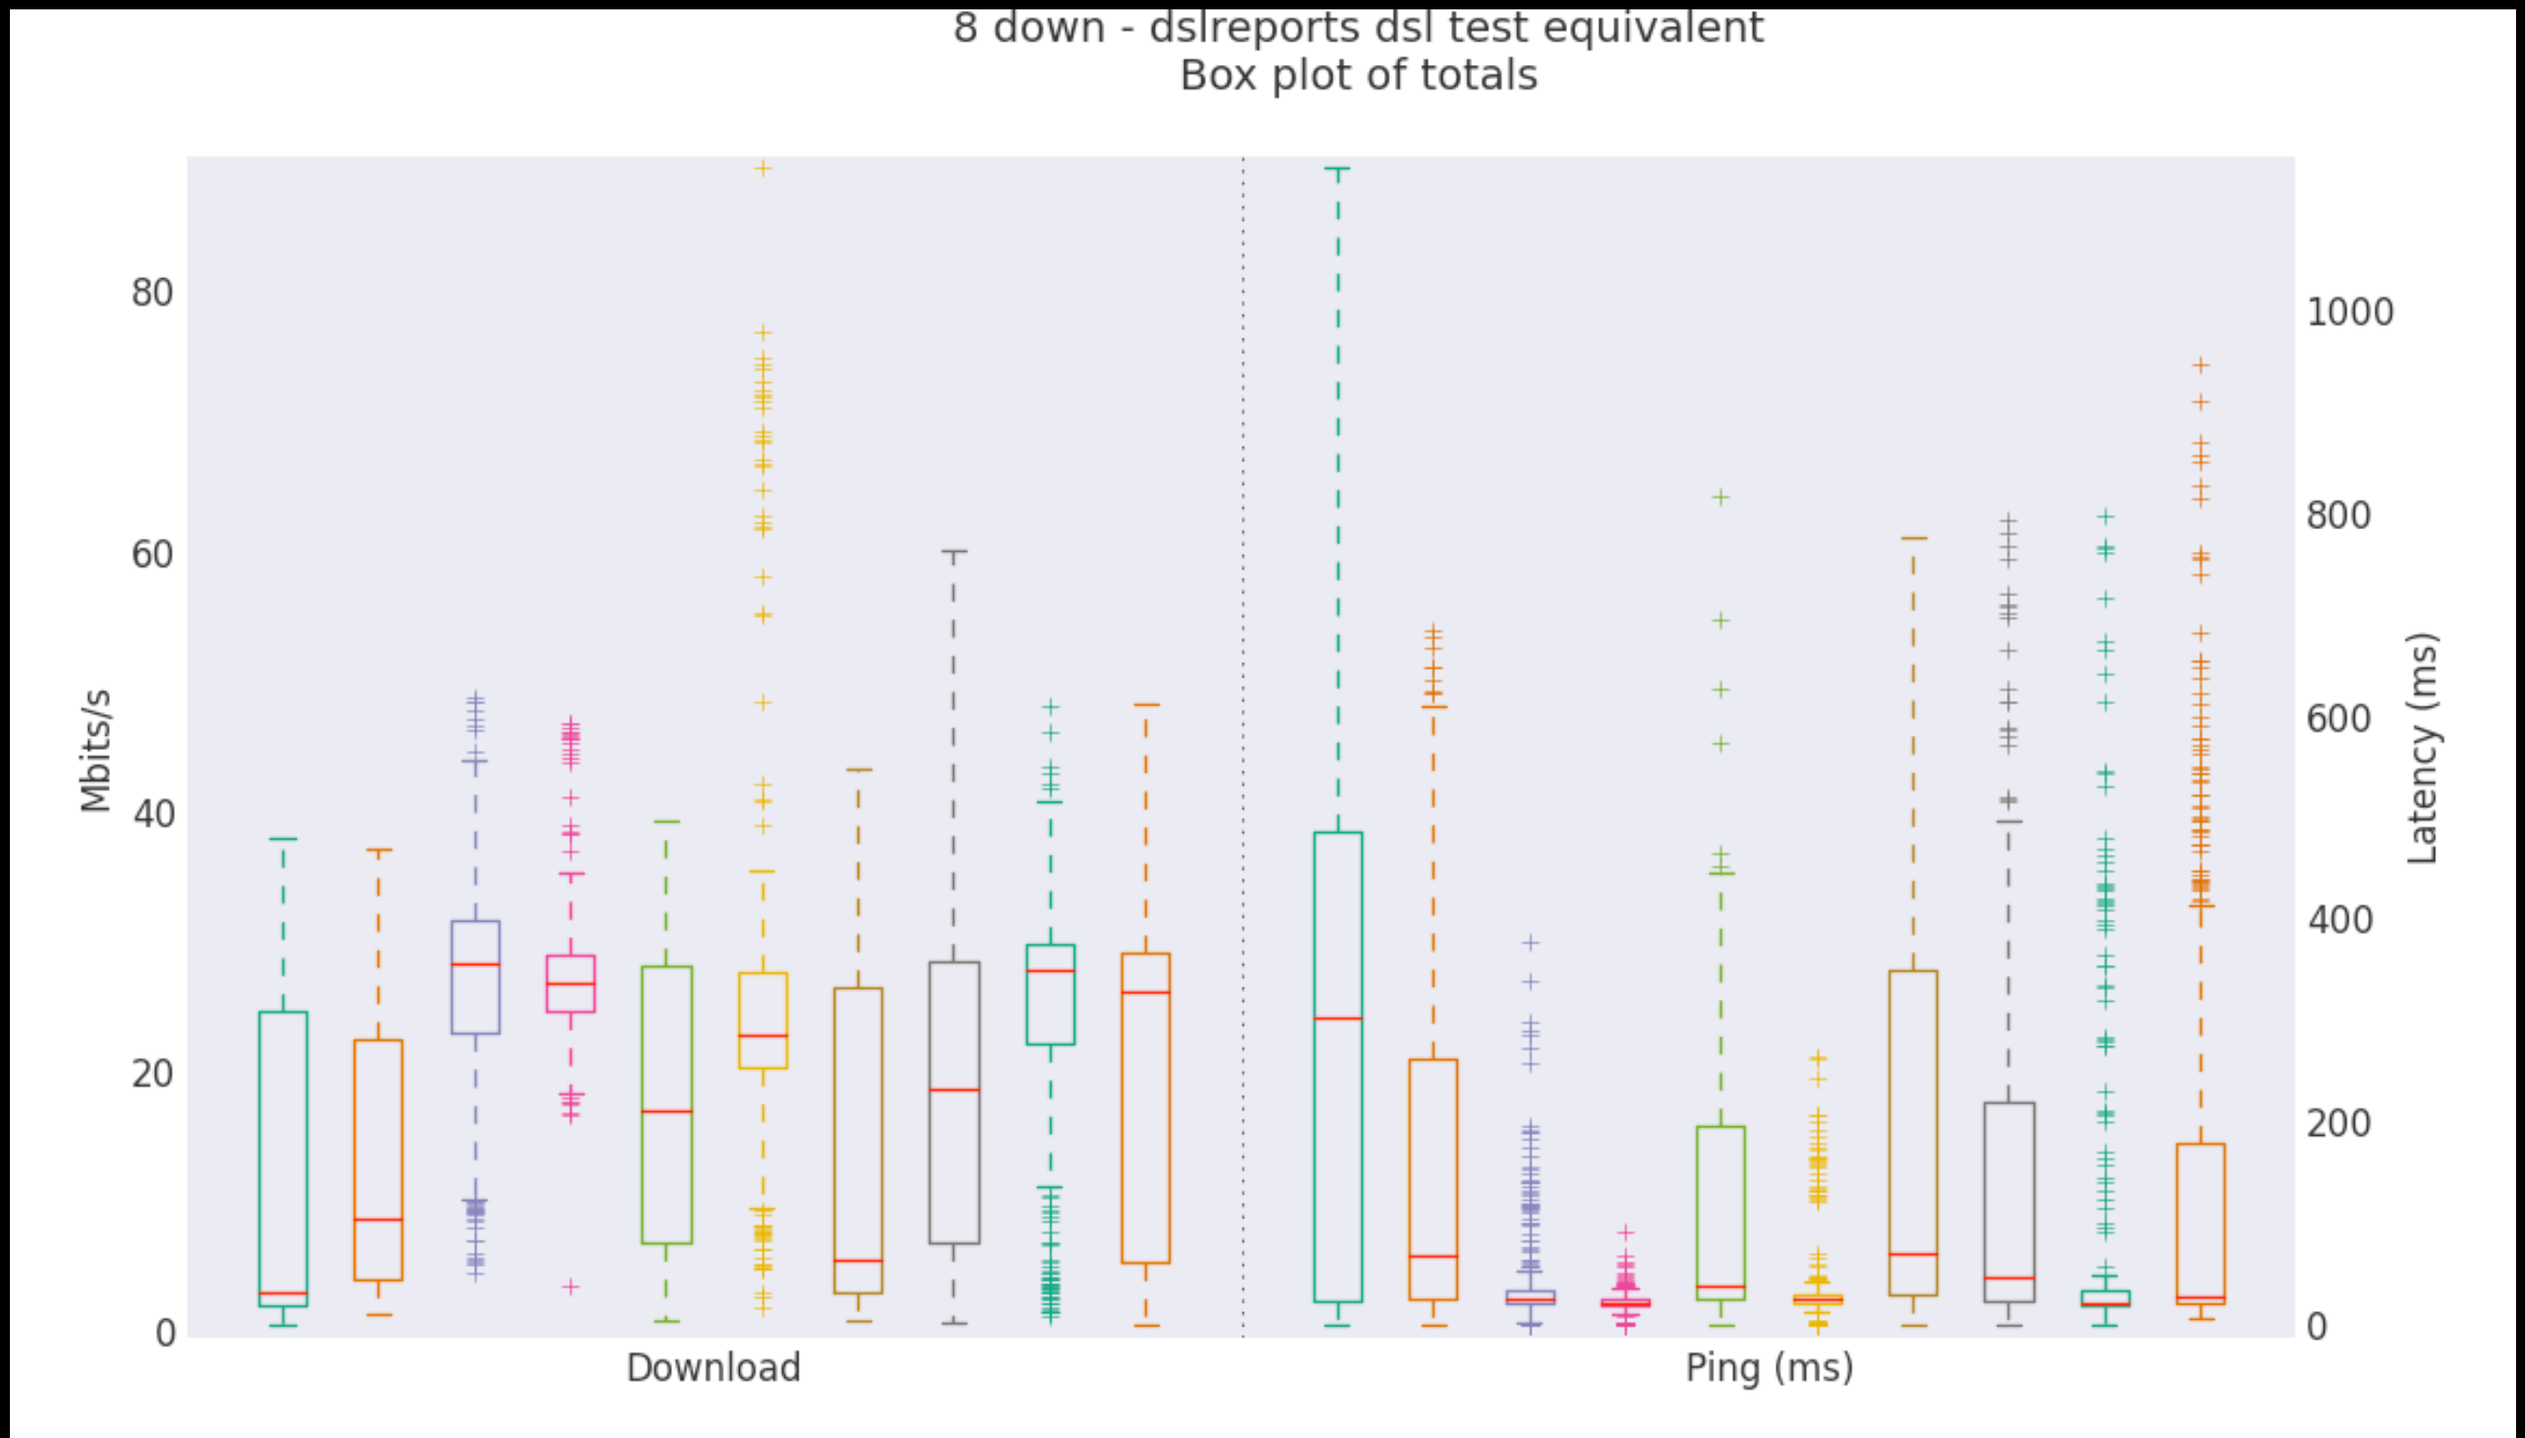

With AQM

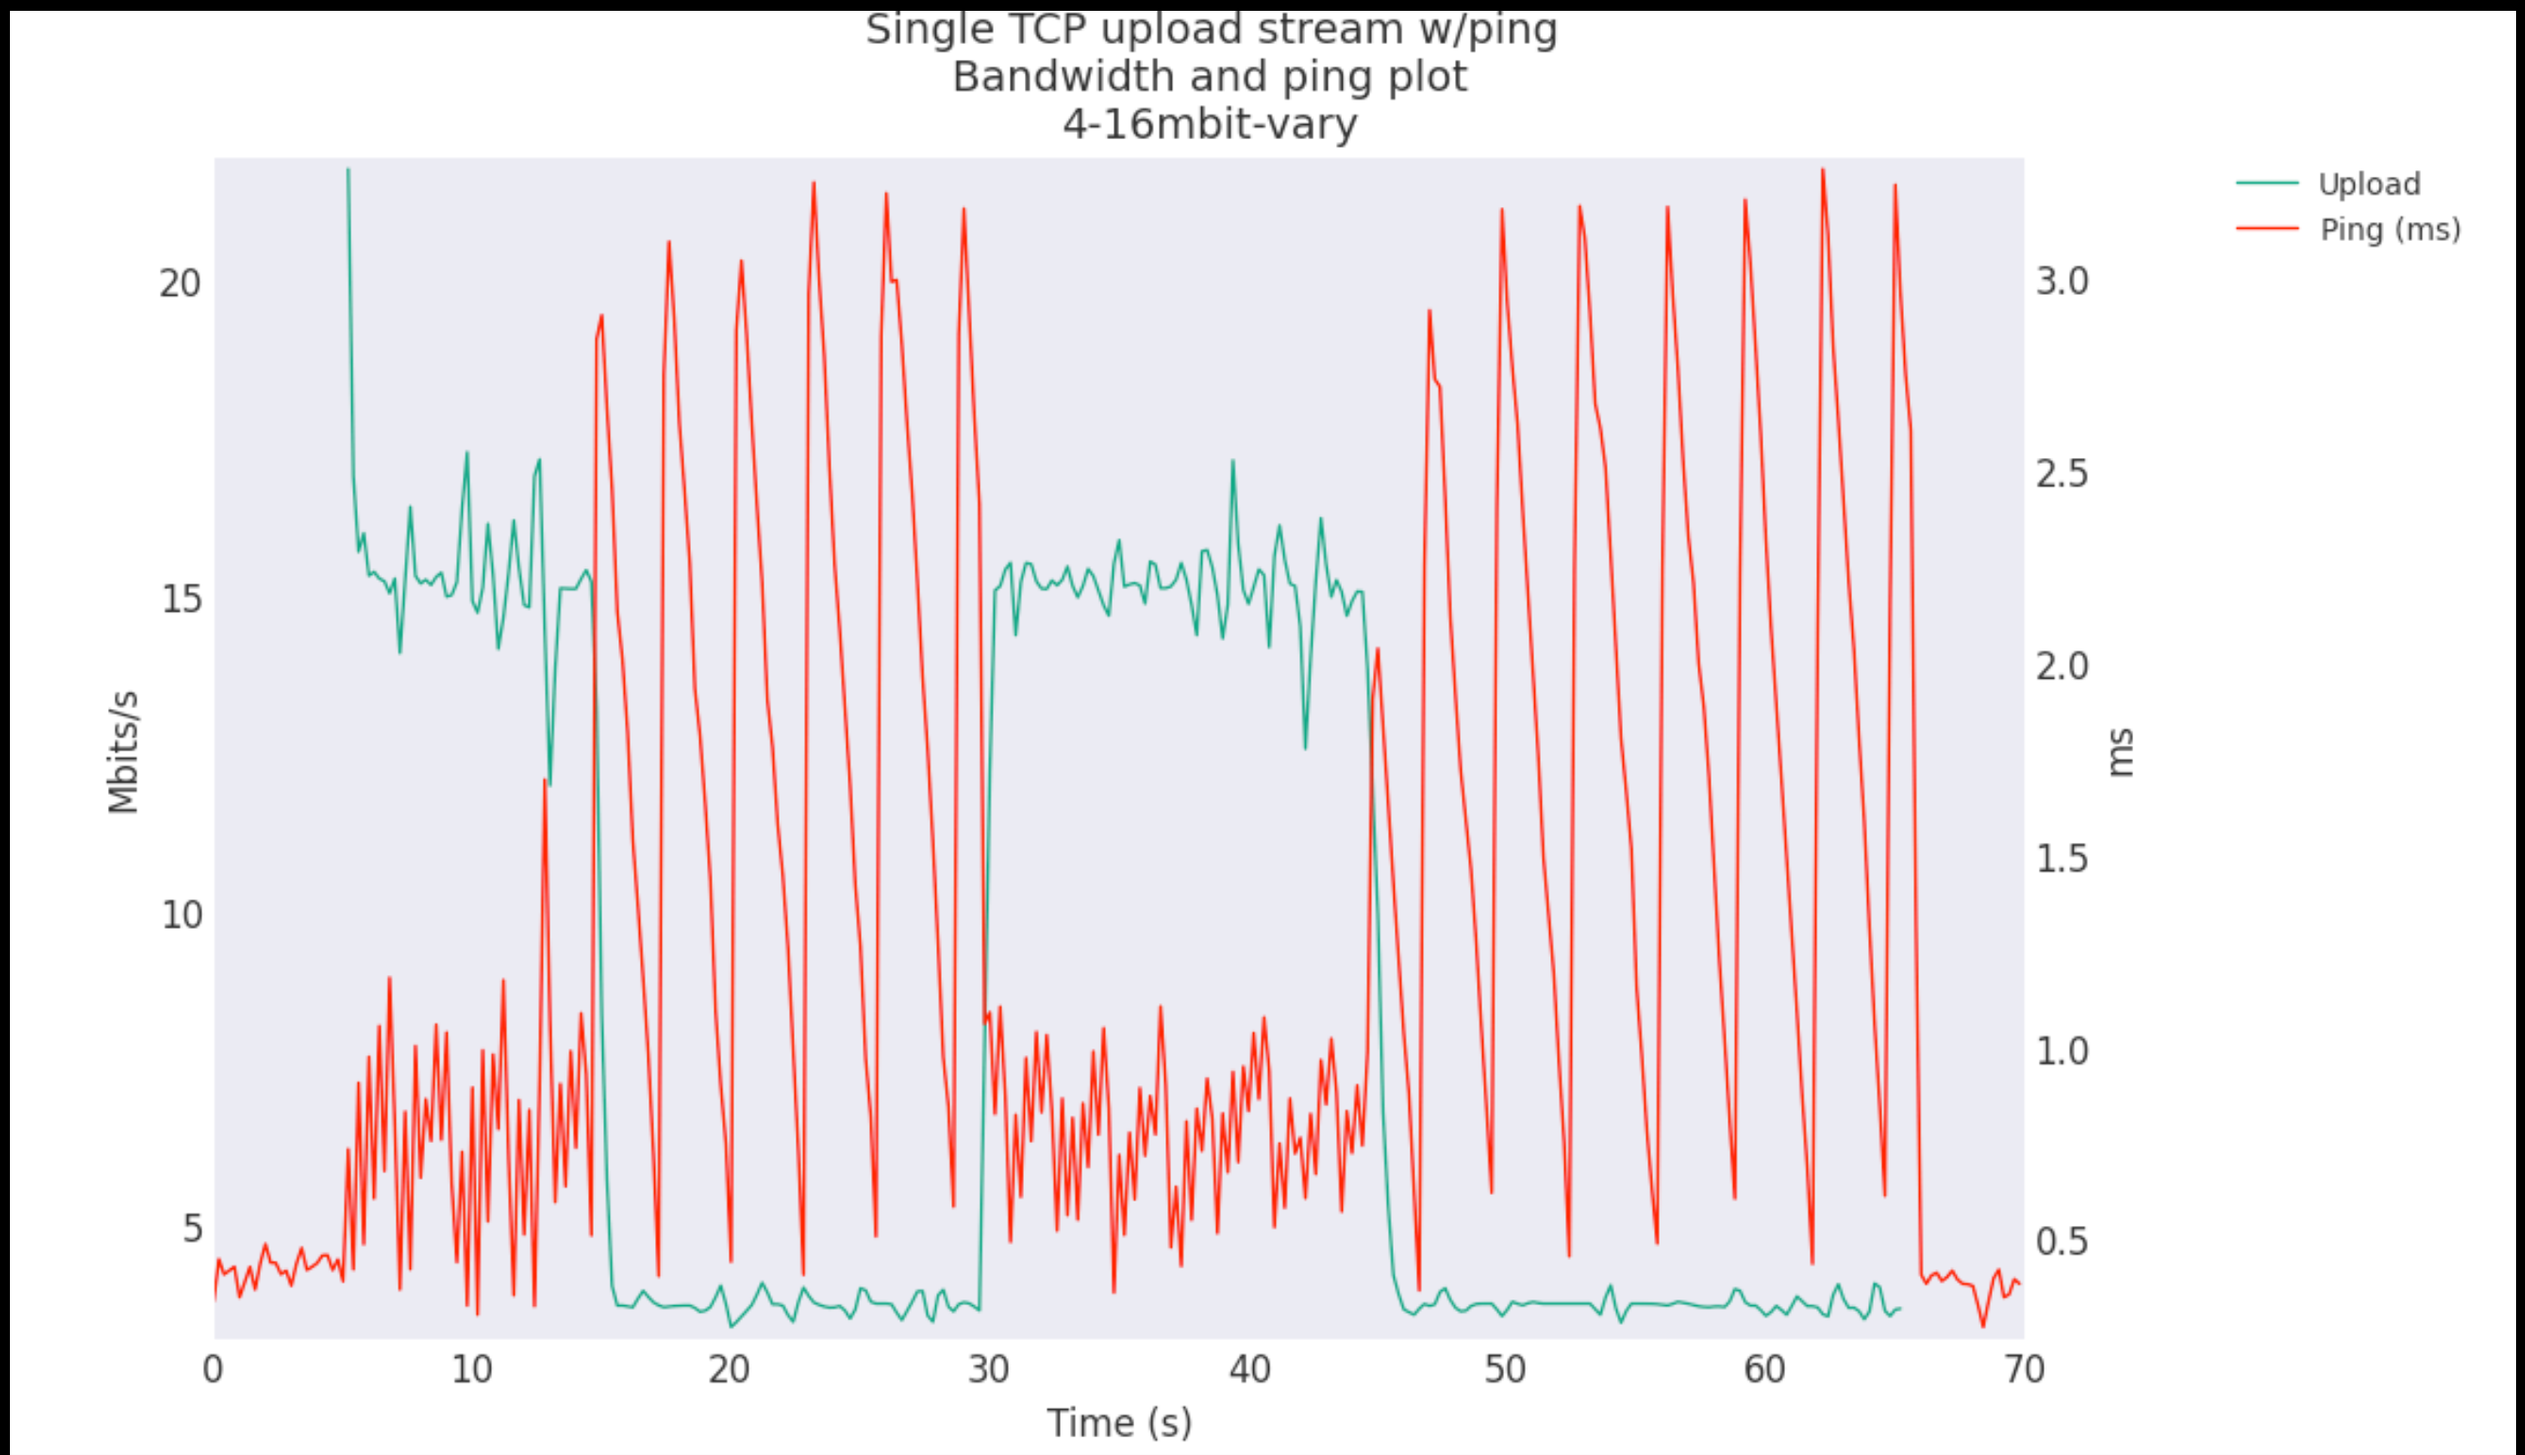
Single TCP upload stream w/ping
Bandwidth and ping plot
4-16mbit-vary
Upload
Ping (ms)
Mbits/s
ms
Time (s)
0
10
20
30
40
50
60
70
5
10
15
20
0.5
1.0
1.5
2.0
2.5
3.0

Flow Isolation

- What most people call "fair queuing".
- Cake began with fq_codel as a core.
- Now has 8-way set-associative hashing.
- Hash collisions virtually eliminated!

Flow Isolation

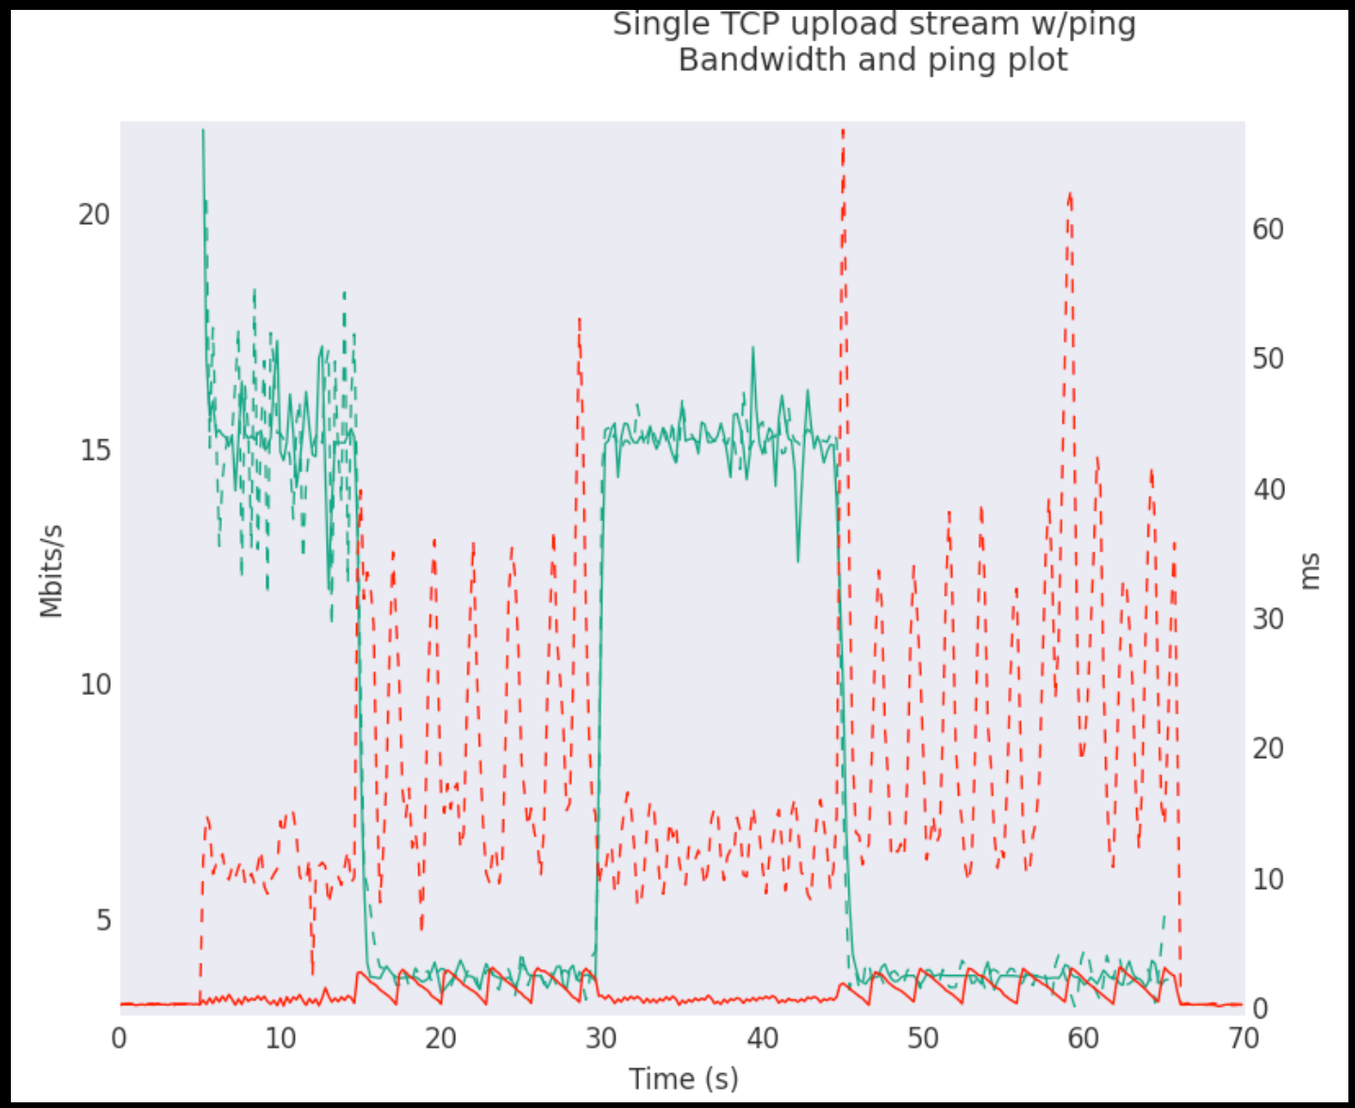

Throughput

Delay (others)

Delay (TCP)

Priority Queuing

- DRR++
  - ...automatically promotes sparse flows.
    - Most latency-sensitive traffic is sparse.
- Four-class Diffserv support...
  - ...without strict priority.
  - ...with soft admission control.

Built-in Shaper

- Take control of the queue!
- Limit channel utilisation per node.
- Works in deficit mode - minimal bursting.
- Tight integration eliminates standing queue between flow-isolator and shaper.

Overhead Compensation

- On wire/air, IP packet encapsulated further:
  - Ethernet frame
  - PPPoE / PPPoA
  - ATM cell quantisation
  - RF pre/postamble
- Cake can account for (some of) these.

Lean & Mean

- 100+ Mbps on a WNDR3700.
- 300+ Mbps on a WRT1900AC v2.
- Less CPU per packet than HTB. Just HTB.
- Less RAM than multiple fq_codels.
- No compromises required.

Archer C7

HTB can’t shape at 115Mbps

Cake can.

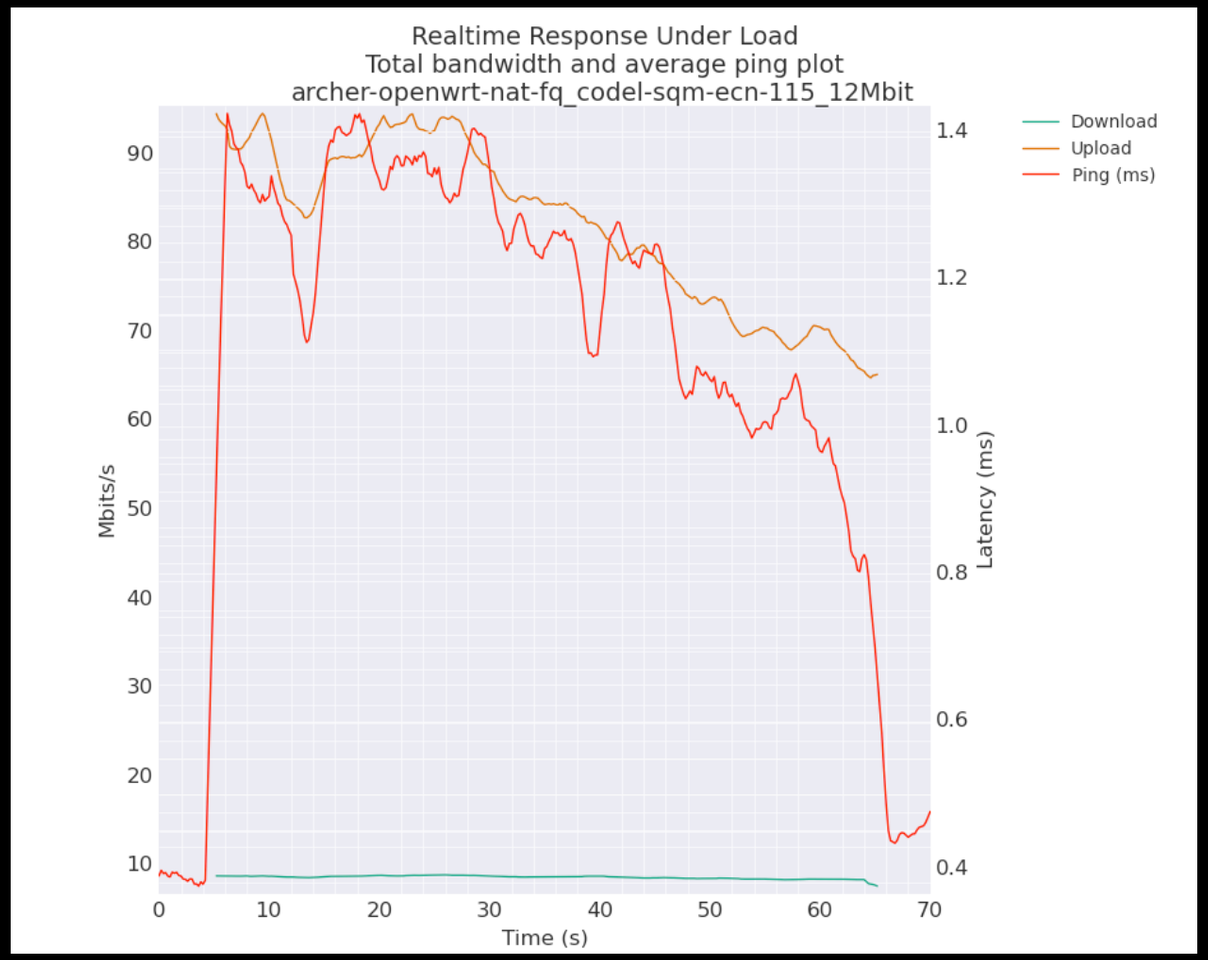

HTB + fq_codel

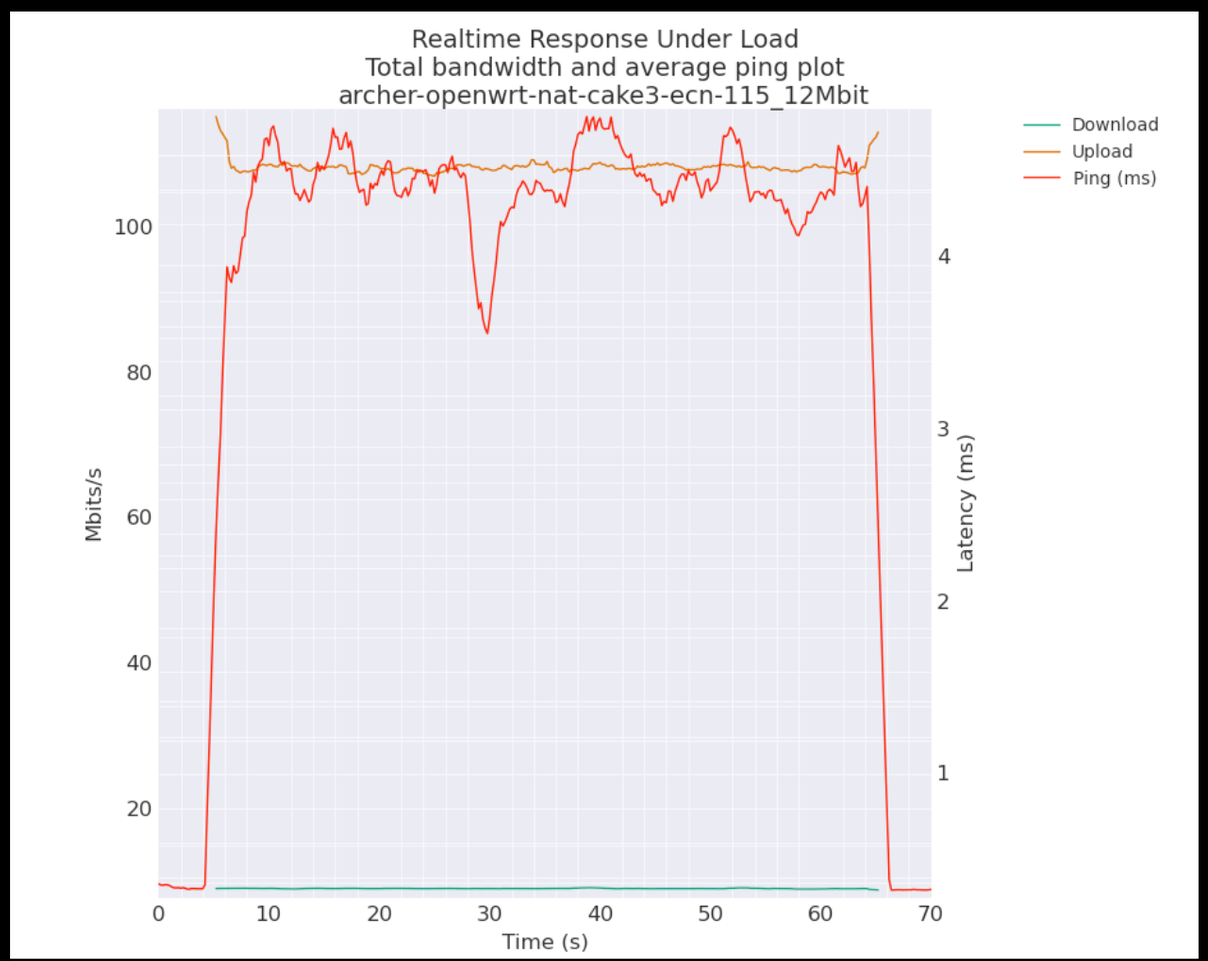

Cake

Lean & Mean

- Replace four 802.11e hardware queues...
  - ...each with their own buffer allocation...
- One queue
  - Less RAM
  - Less latency
  - No strict priority - no starvation!

Easy to Configure

- One stop shop - one ‘tc’ invocation per interface.
- Sensible defaults for unspecified params.

```
tc qdisc add dev eth0 root handle 1: cake bandwidth 10Mbit
```

- Change parameters on the fly:

```
tc qdisc change dev eth0 handle 1: cake ...
```

Without losing packets!

Modelling Bandwidth Changes

- Change cake's bandwidth on the fly...
  - ...model a variable Wi-Fi link!

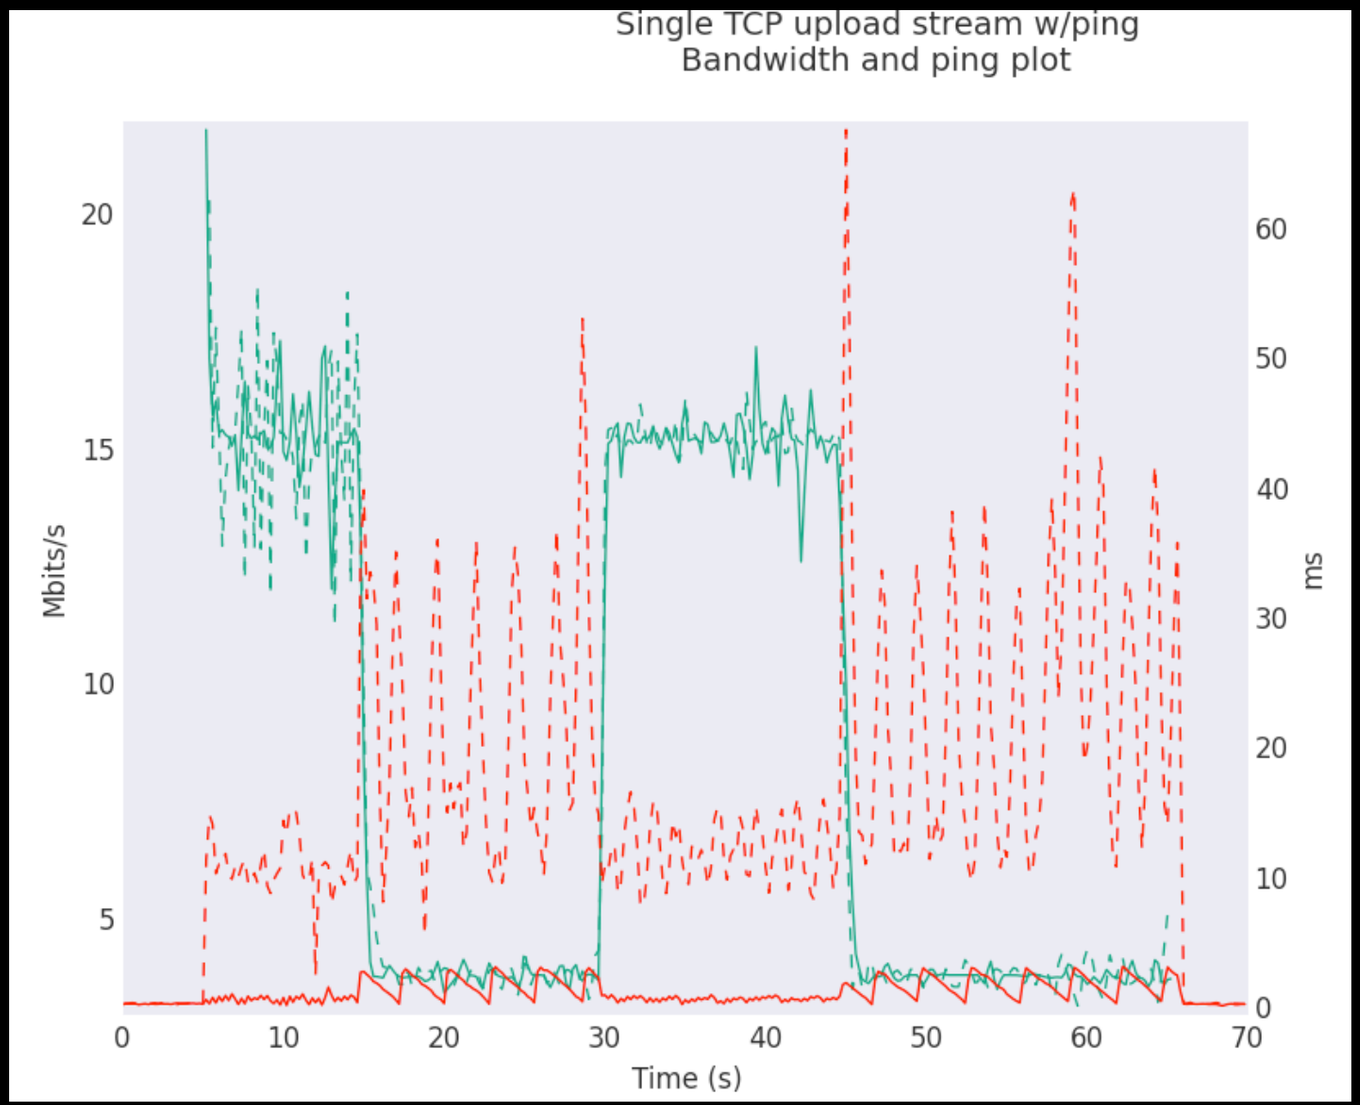

Throughput

Delay (others)

Delay (TCP)

Easy to Configure

- One stop shop...
- Concise shortcuts for common ISP framings:

```
pppoa-vcmux   -> atm overhead 10     # Efficient ADSL
pppoe-llcsnap -> atm overhead 40     # Lazy ISP's ADSL
bridged-ptm   -> noatm overhead 19   # VDSL
```

...and more...

Summary

- Efficient shaper (replaces inefficient HTB)
- Diffserv support (which is halfway sane)
- Does everything fq_codel does (but better)
- Eating my own dogfood!
- Still being improved...

How?

- Get out-of-tree kernel module

```
git clone https://github.com/dtaht/sch_cake
```

- Get patched iproute2

```
git clone https://kau.toke.dk/git/cake/iproute2
```

- Get OpenWRT packages

```
git clone https://github.com/dtaht/ceropackages-3.10
        cd net/kmod-sched-cake
        cd net/tc-adv
```