



What is Bufferbloat?



Stephen Hemminger
stephen@networkplumber.org



Solving Plumbing problems

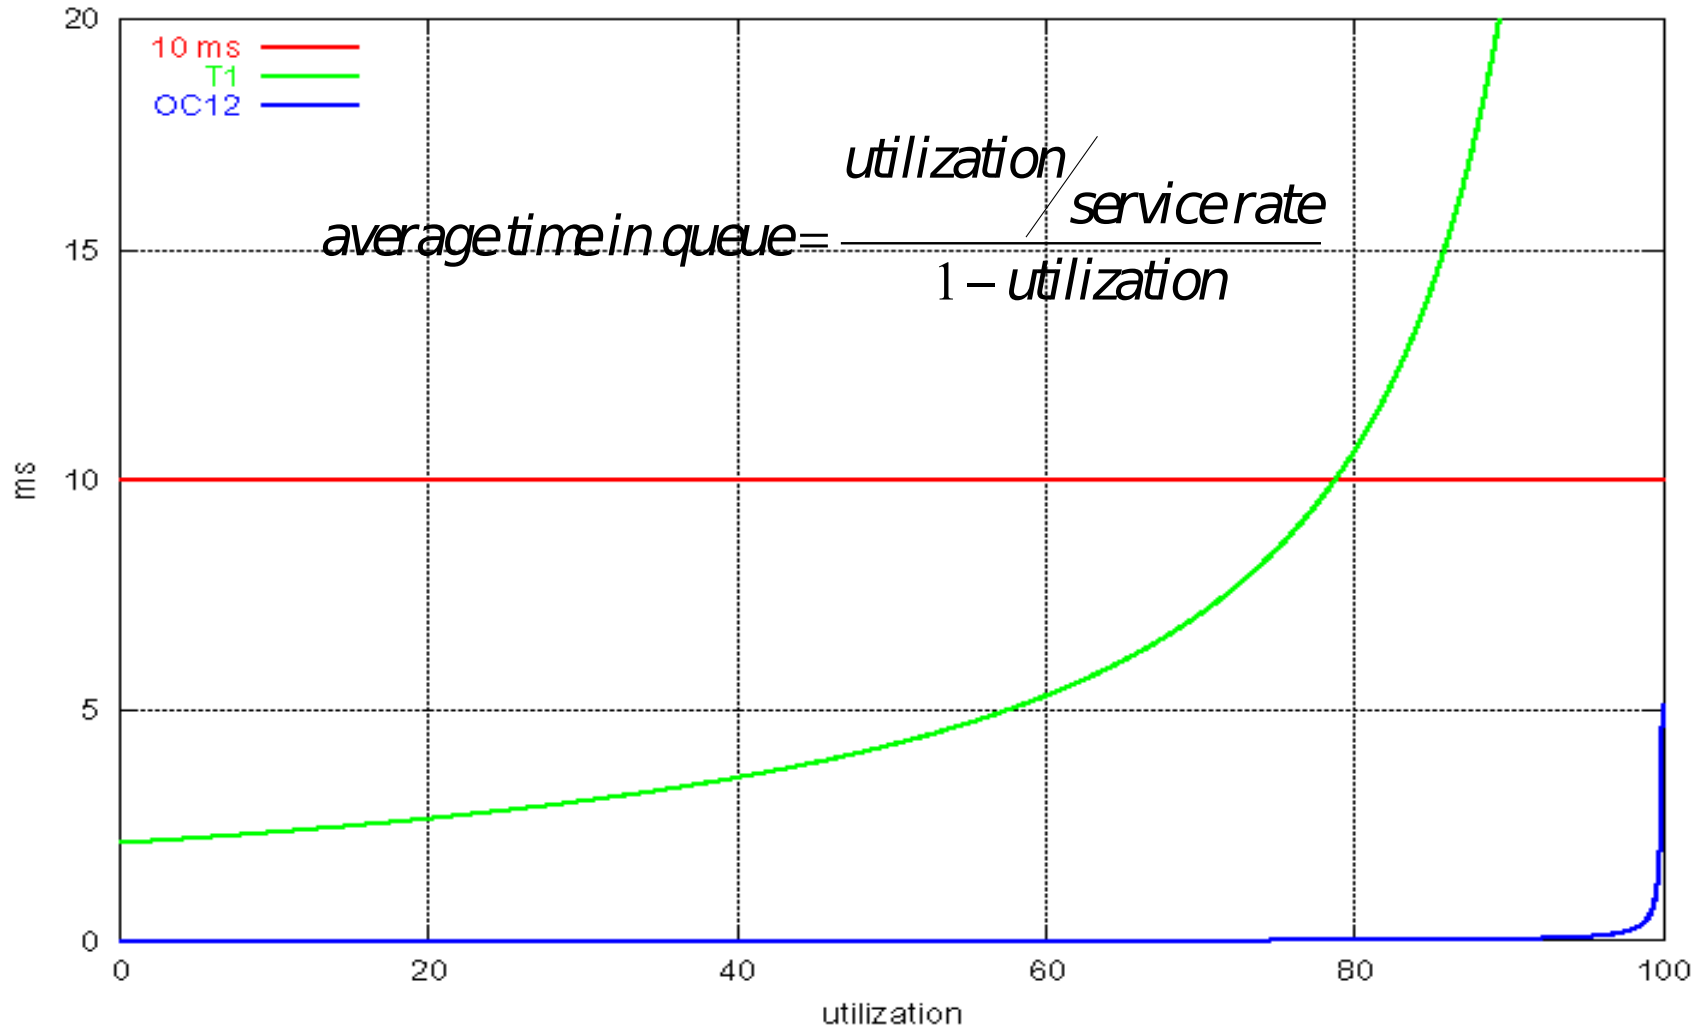
- Discovery
- Blame
- Enlightenment
- Nirvana



Bufferbloat: definition



Queuing Theory Basics



From Fred Baker: Bufferbloat!
Graphic courtesy Sprint, Apricot 2004

How bad is it?

150ms – VOIP

20ms – Desktop lag

100ms – DNS lookup fail

ARP, ND – Can't find machine

1000ms – User's give up


scp copy fails

DHCP – Machine falls off network

Triggers


- Uploads

- Youtube
- Dropbox
- Torrent
- Large Email attachments
- Backup services

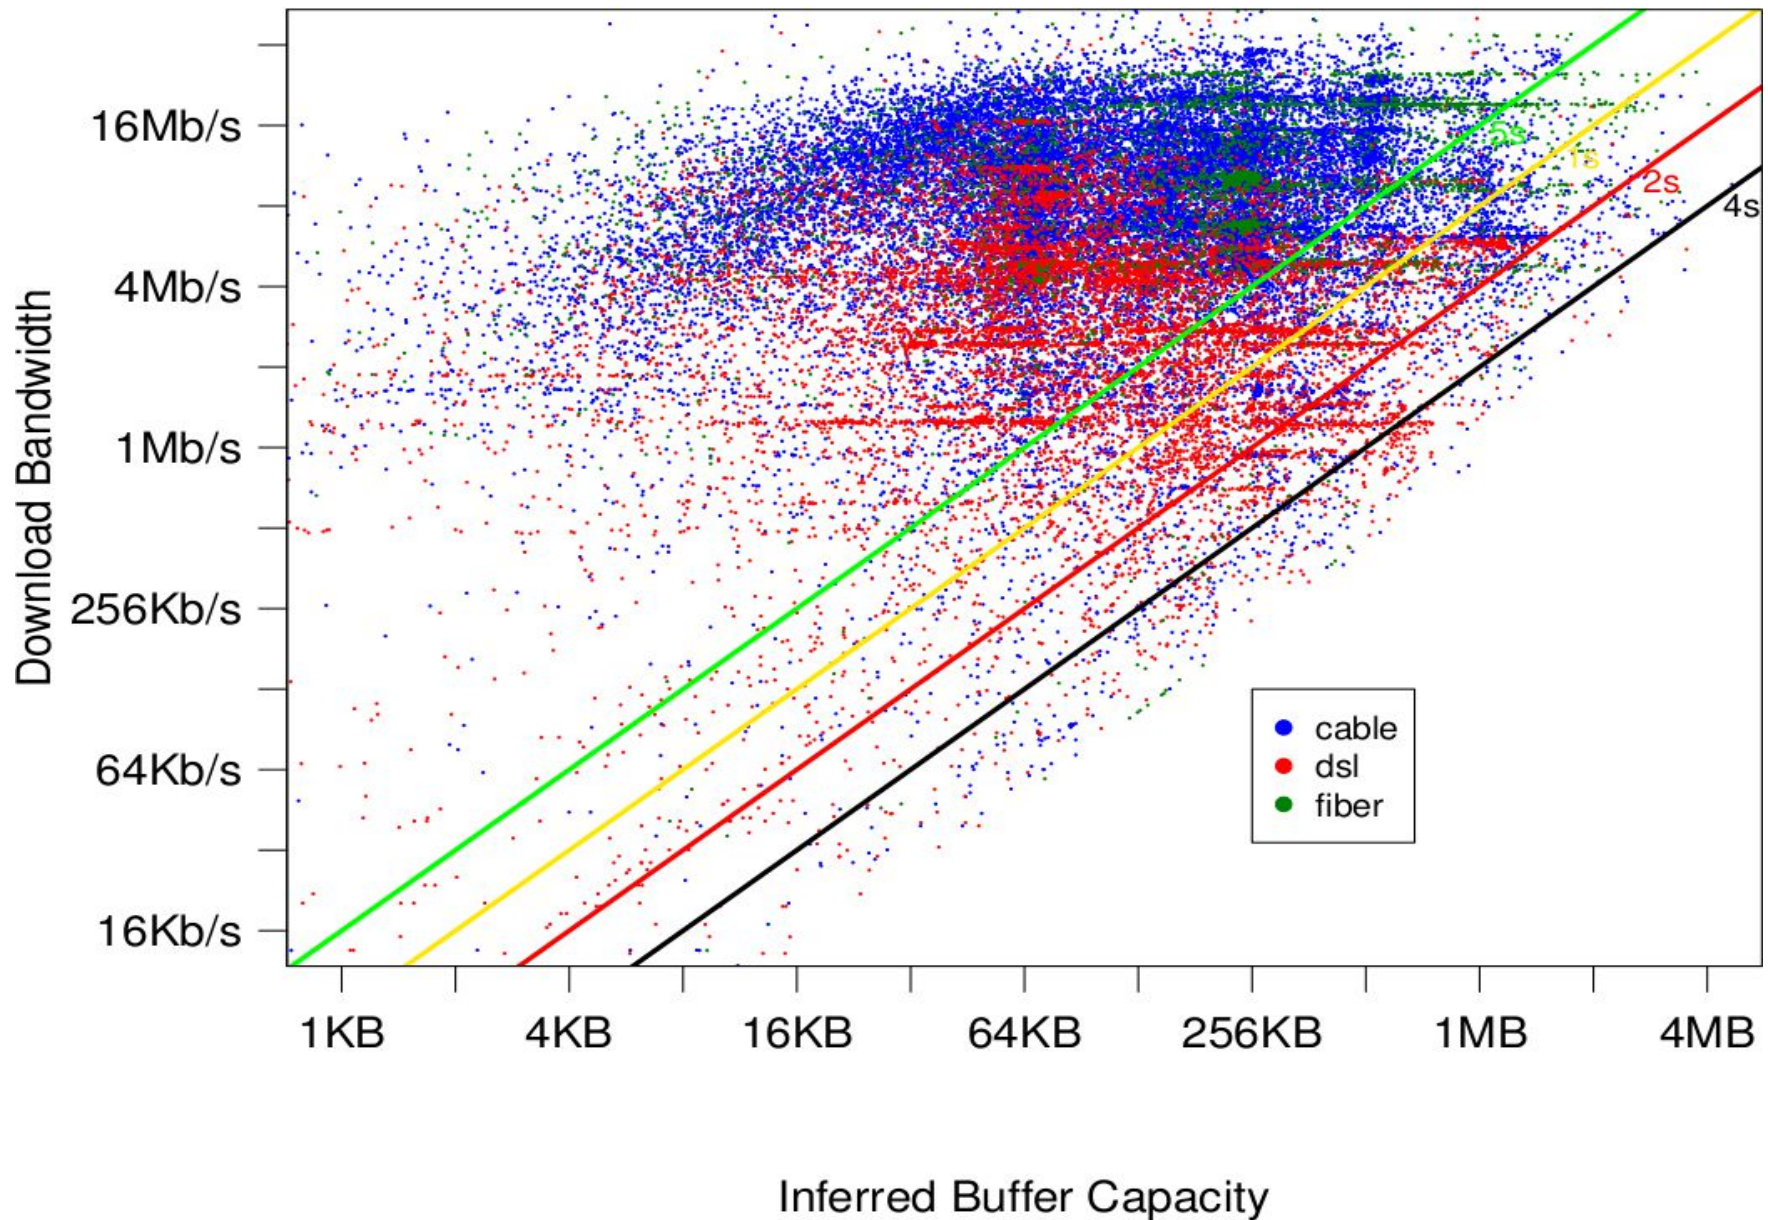


- Downloads

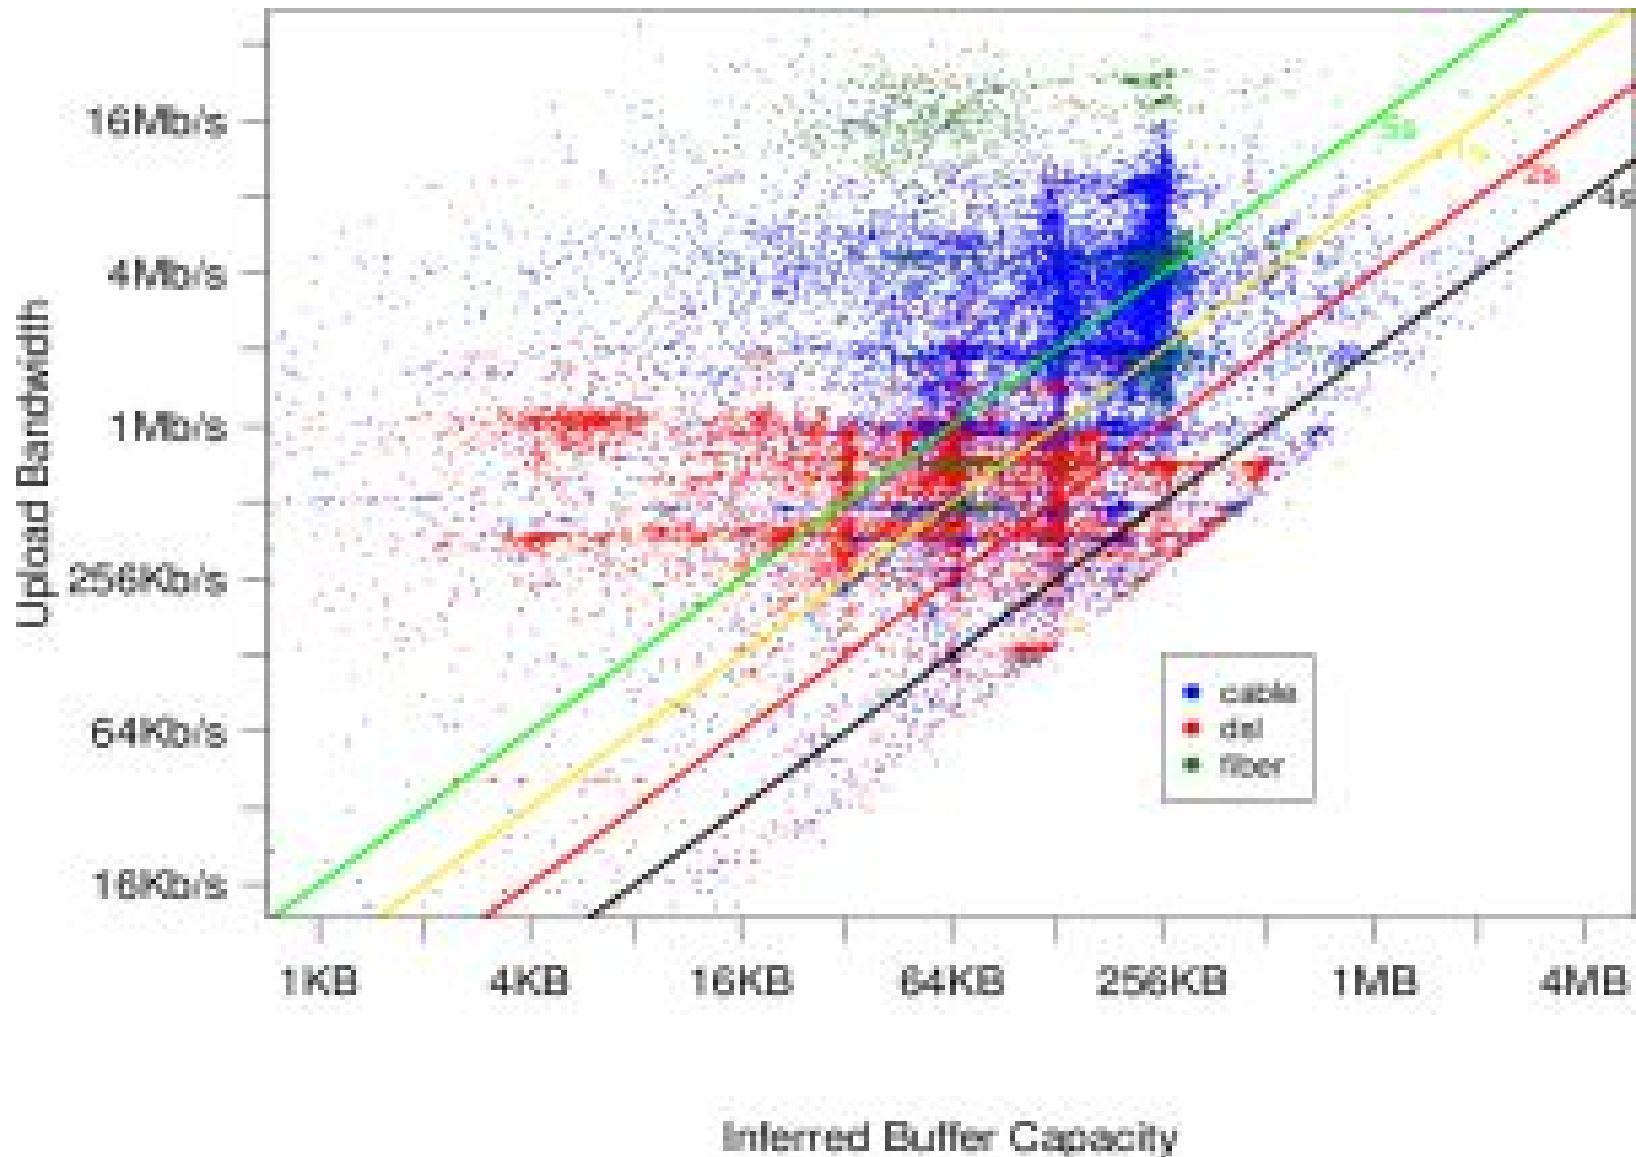
- Movie downloads
- Image heavy websites
- Video teleconferencing



Netalyzr: Downstream

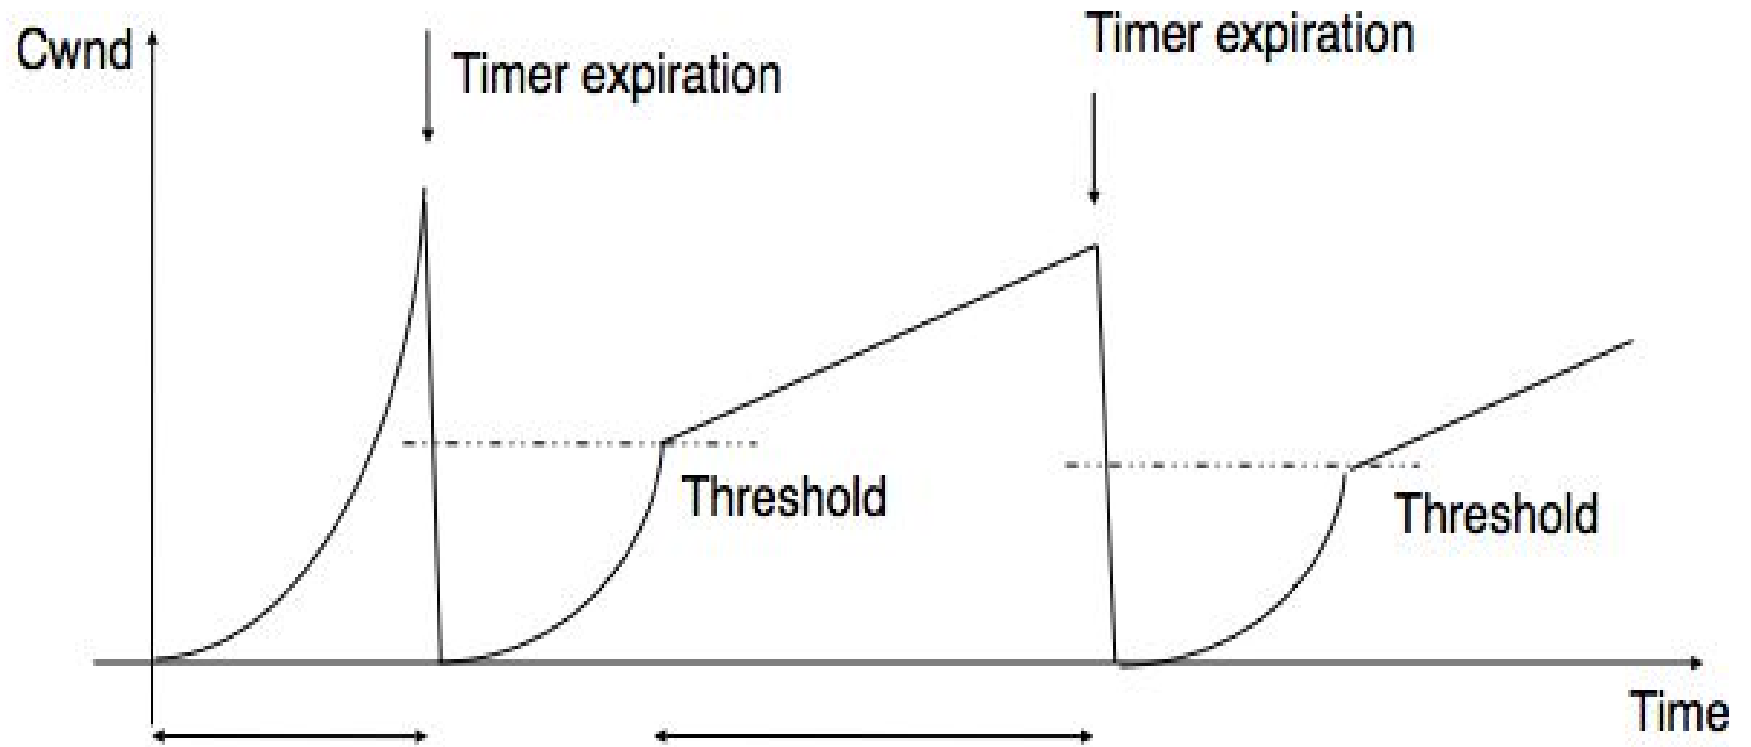


Netalyzr: Upstream




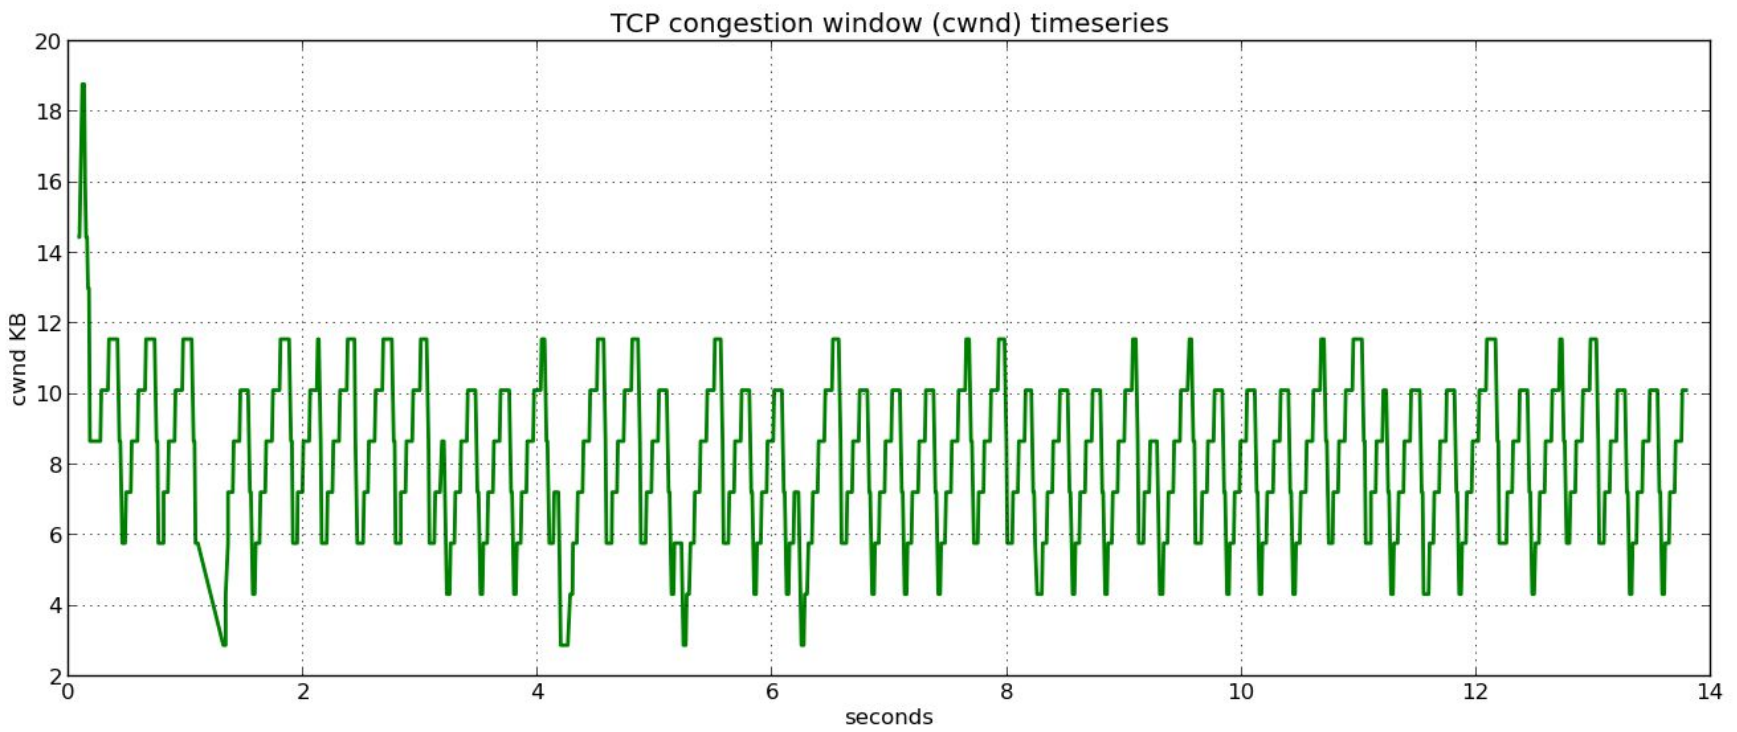


Classical TCP




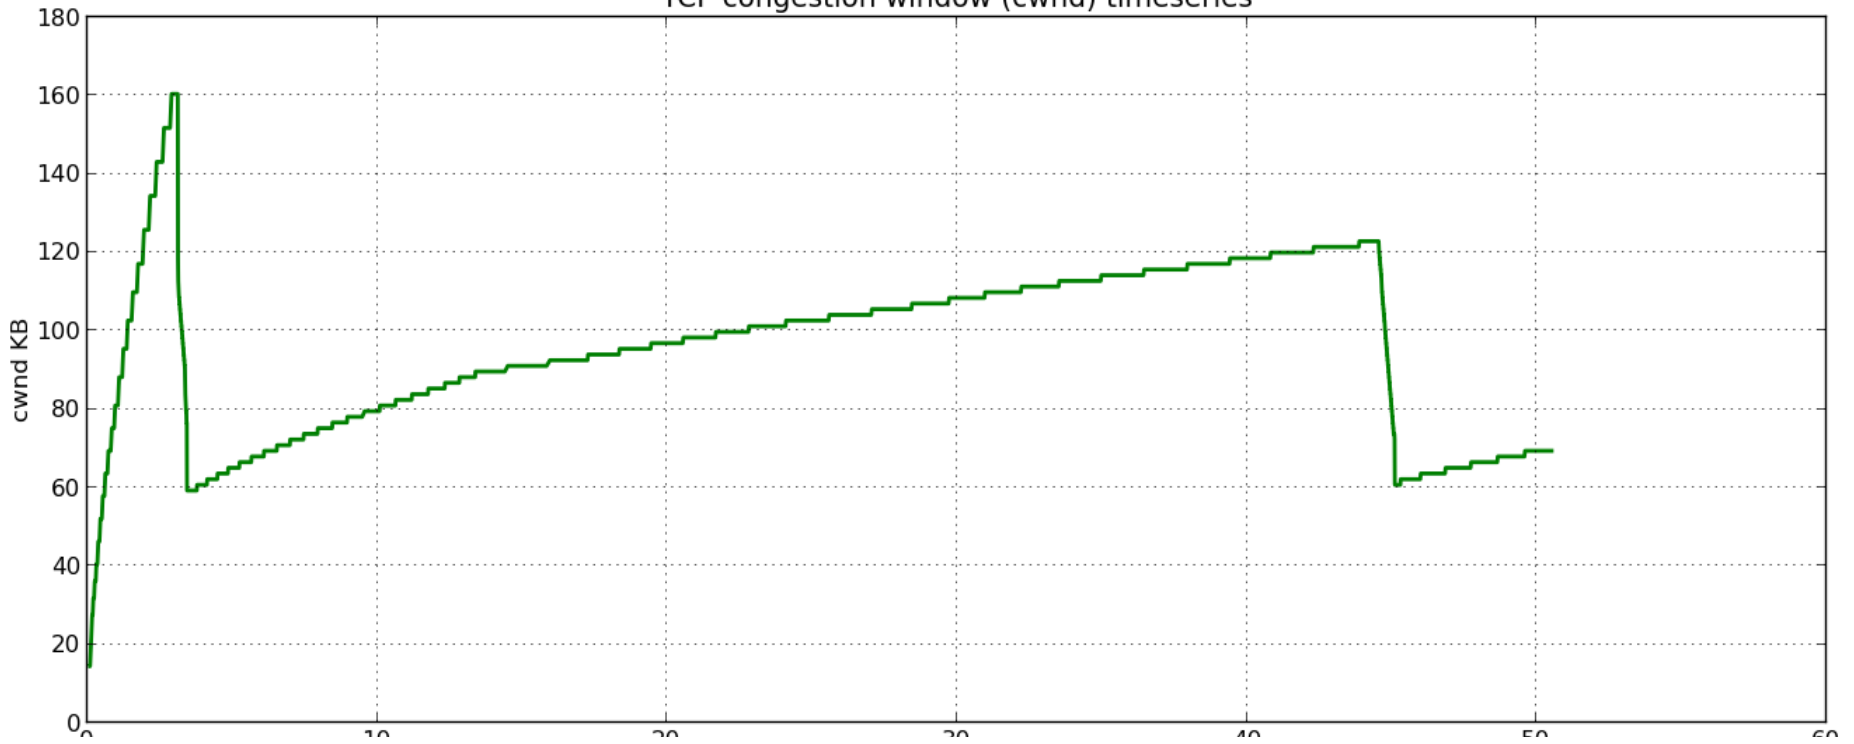
Slow-start
exponential increase of cwnd

Congestion avoidance
linear increase of cwnd




Antonin Bas
Stanford Cs244 assignment


TCP congestion window (cwnd) timeseries



Two TCP CUBIC flows, 1GB transfer, 100 Mbit/sec RTT=100ms Q=488KB



TCP throughput dynamics



From Fred Baker: Bufferbloat!

Blame Linux



Windows XP

- Maximum window 64k



Windows 7

- Bandwidth limit to 80 mbits




Android

- Receive window limited

Blame the customer

- Customers call support
- Applications are using more bandwidth
- Block and charge






Back to Reality

Benchmark issues

- Bad
 - Bytes/sec
 - Packes/sec
 - Latency

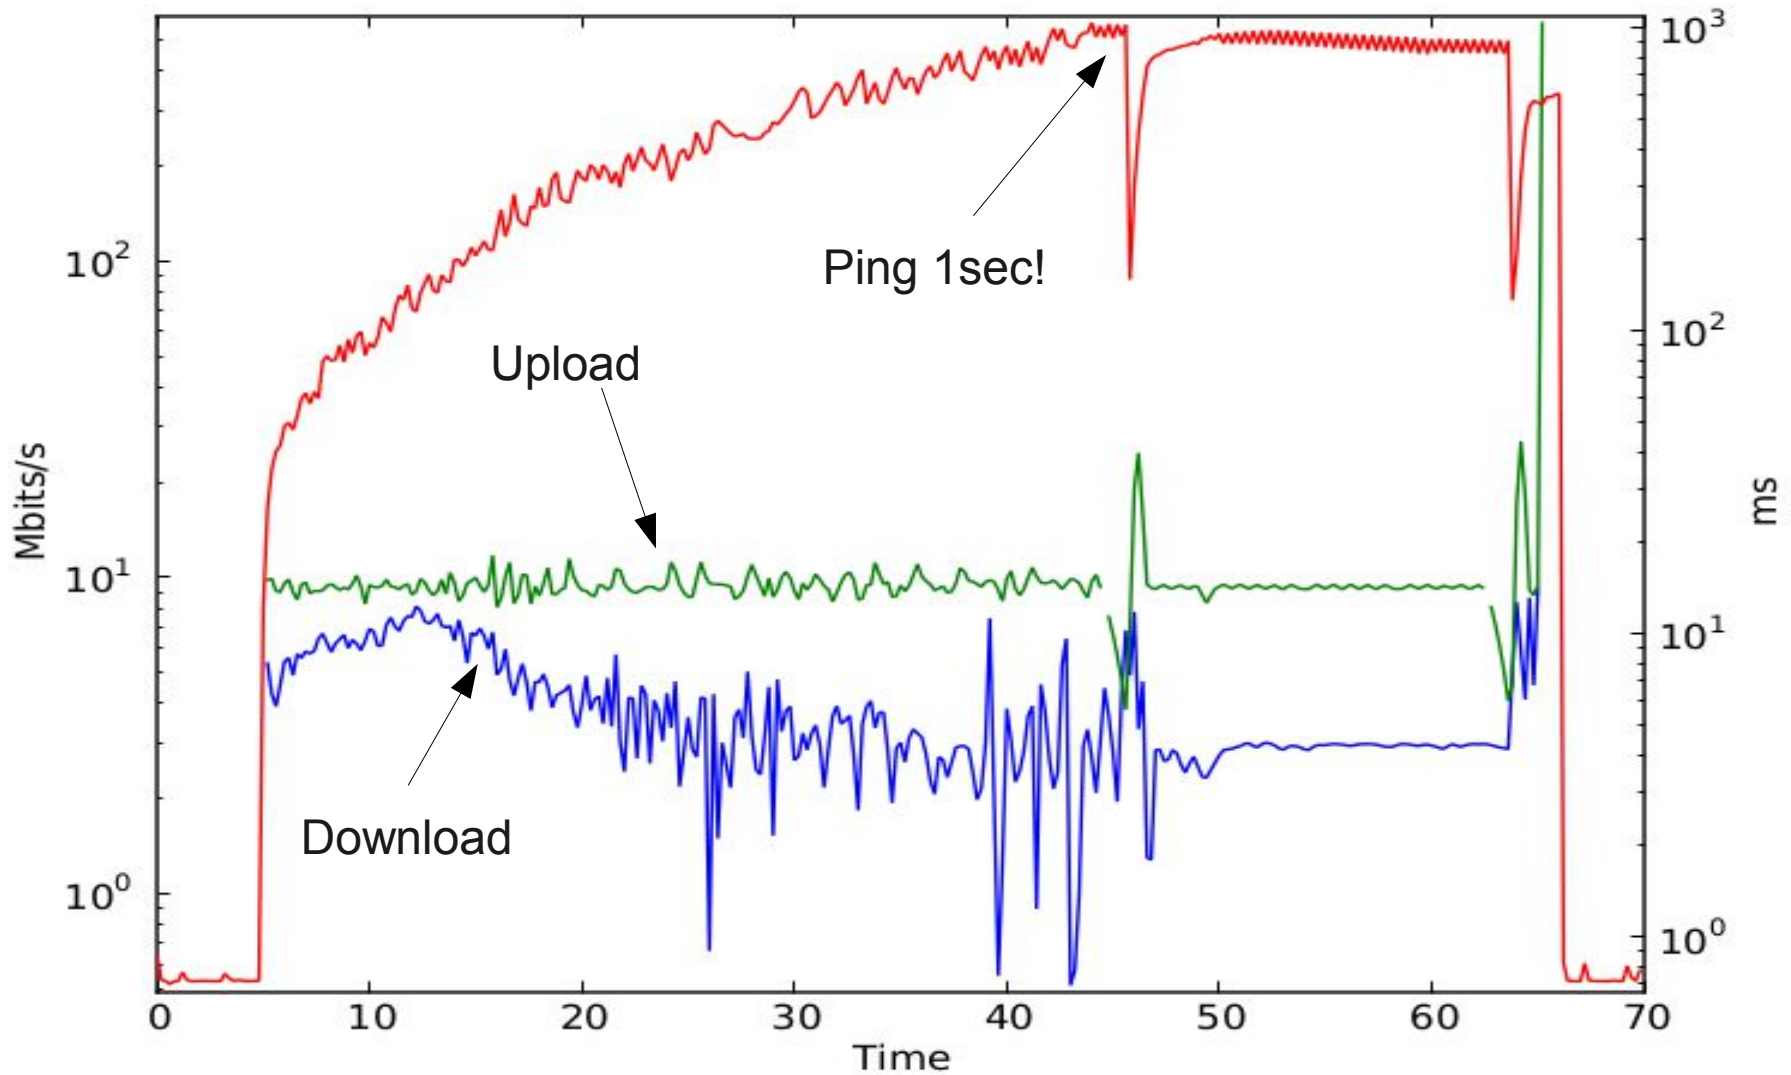


Making a good test


- Combinations
 - RRUL




FIFO – the default




Stochastic Fair Queue



Controlled Delay - codel



Fair Queue Controlled Delay



Fight

Bufferbloat



smetek.de

Linux Everywhere



In the Datacenter

In the Home



If we fix Bufferbloat, the world will change... eventually


The ICSI Netalyzr

Network Access Link Properties + -

- Network performance (?): Latency: 220 ms, Loss: 0.0%** +
- TCP connection setup latency (?): 229ms** +
- Background measurement of network health (?): no transient outages** +
- Network bandwidth (?): Upload 3.8 Mbit/s, Download 250 Kbit/s** -
- Your Uplink** We measured your uplink's sending bandwidth at 3.8 Mbit/s. This level of bandwidth works well for many users.
- Your Downlink** We measured your downlink's receiving bandwidth at 250 Kbit/s. This rate could be considered quite slow, and will affect your user experience if you perform large transfers.
- Network buffer measurements (?): Uplink is good, Downlink 490 ms** -
- We were not able to produce enough traffic to load the uplink buffer, or the uplink buffer is particularly small. You probably have excellent behavior when uploading files and attempting to do other tasks.
- We estimate your downlink as having 490 ms of buffering. This level can in some situations prove somewhat high, and you may experience degraded performance when performing interactive tasks such as web surfing while simultaneously conducting large downloads. Real time applications, such as games or audio chat, may also work poorly when conducting large downloads at the same time.

Bufferbloat project

- Information
 - Research
 - Presentations
 - Mailing lists



Status System Services **Network** Logout

Interfaces Wifi Switch DHCP and DNS Hostnames Static Routes Diagnostics Firewall Radvd Wake on LAN
 AHCP Server **QoS**

Quality of Service

With QoS you can prioritize network traffic selected by addresses, ports or services.

Interfaces





GE00  Delete


Enable	<input checked="" type="checkbox"/>
Classification group	default
Calculate overhead	<input type="checkbox"/>
Half-duplex	<input type="checkbox"/>
Download speed (kbit/s)	3000
Upload speed (kbit/s)	750




1. Uncheck "Enable" to turn off QoS
2. Use <http://speedtest.net>
3. Enter Download and Upload speeds
4. Check "Enable" to turn QoS back on
5. Click Save & Apply

 Add

Classification Rules

Target	Source host	Destination host	Service	Protocol	Ports	Number of bytes	Sort
prior	all	all	all	TCP	22,222,53	<input type="text"/>	↑ ↓ 
norm	all	all	all	TCP	20,21,25,80,81,8123,110,4	<input type="text"/>	↑ ↓ 
expr	all	all	all	UDP	53,123,5060,5190	<input type="text"/>	↑ ↓ 
low	all	all	all	all	873	<input type="text"/>	↑ ↓ 

 Add

 Reset  Save  Save & Apply

Future Research

- Systemic view
- Change Linux default?
- Wireless
- High speed 10/40/100G?

Thank you

Stephen Hemminger
stephen@networkplumber.org