

# **Bufferbloat Research Tpics**



Jim Gettys

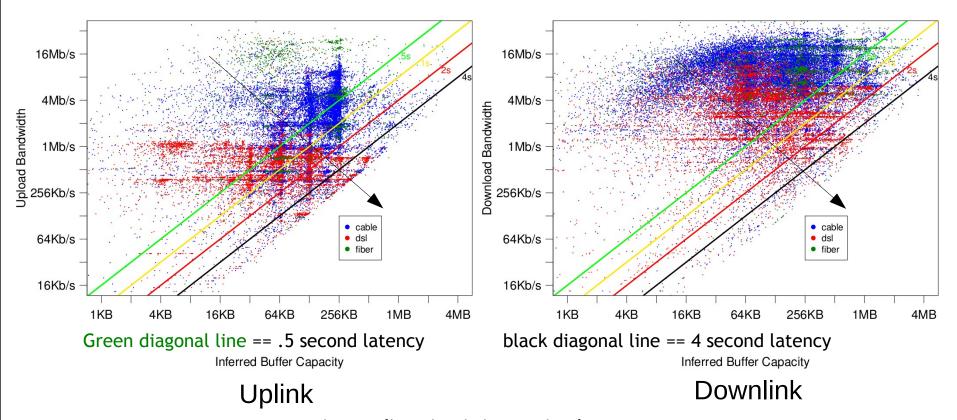
**Bell Labs** 

July 29, 2012



james.gettys@alcatel-lucent.com, jg@freedesktop.org

#### "Netalyzr: Illuminating Edge Network Neutrality, Security, and Performan C. Kreibich, N. Weaver, B. Nechaev, and V. Paxson



Arrow direction is increasing latency
Note: telephony standards for latency are maximum of 150ms!!!

This data is a *lower* bound on the severity of the broadband bufferbloat problem.

# "Global" Research Topics



Is bufferbloat getting better? Or worse? And where?

- Can you exploit NTP to do passive measurements?
- Active measurements? Mlabs data analysis? FCC analysis?

Which links and networks are broken?

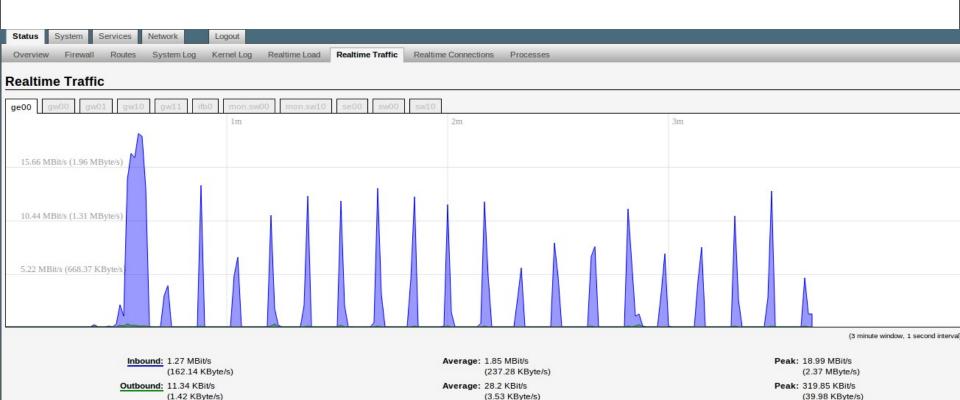
Are there techniques to find "dark" buffers?

Are new applications causing unnecessary pain? (e.g. current stupid video "streaming" players could be better)

What are better strategies for those applications that inflict less harm while the network is being fixed?

# Netflix "streaming" Behaviour





# Algorithmic Questions



Other adaptive AQM algorithms certainly exist. Invent a better one...

Is CoDel algorithms suitable in the data center? If so, what changes are necessary to its "constants"?

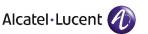
How does packet aggregation (e.g. 802.11n, DOCSIS) affect AQM algorithms?

You have many more cycles/packet available at the edge of the Internet, and bandwidth is always hard to come by there. What are smarter queuing algorithms, most suitable for different locations in the network, since fq\_codel demonstrates that smarter queuing is feasible on today's fast processors that were infeasible 20 years ago?

Your bandwidth suddenly dropped by a factor of 10. What is the best strategy to bring a queue back down quickly?

Wireless bandwidth is varying at a faster timescale than the AQM can handle. Now what? What's the best strategy?

Exactly how should we do ECN marking? When do we drop?



### Data Acquisition & Analysis



What is going on on a busy 802.11 network?

How broken are cellular networks? Where?

Existing traces of busy networks are often completely misleading due to bufferbloat? How will things change as we fix bufferbloat?

Data will need to inform the design of smarter queuing strategies. How do we collect such data given experiments and get it into the hands of people working on such algorithms?

Testbeds for wireless are few and far between for most researchers and developers. So issues of scale will be missed. How do they get access to supported testbeds? Do we have enough of them?

#### **Tools**



Basic tools are often unavailable or unmaintained. Example:

- Pathchar source never released. Binaries have rotted
- Pchar has no active maintainer

Graphical tools for bufferbloat may make it much more understandable to journeyman network engineers; ping, traceroute are insufficient

How do you best identify bloated bottlenecks? How do you detect these transients of bad behavior, so they get fixed before they are a crisis?

ECN testing needs to continue: studies such as Bauer, Beverly et. Al need to continue, along with debugging tools

Consumer bufferbloat tests are needed: without them, progress will be slowed and unnecessary finger pointing will be amplified

#### Home Router Disaster



#### Commercial home routers are broken in 4 major ways

- Firmware is horribly antique and insecure; today's latest commercial home routers usually ships (at least) 5 year old software on new hardware, which seldom if ever is updated once "stable", which then rots for years after that without update
- Decent IPv6 deployment is now gated by the home routers
- Extreme bufferbloat in all its forms
- Tragedy of the Commons: Funding model of the home router market is broken; there is next to no funding toward engineering to fix problems today: this means that little will happen without community participation

Time to roll up your sleeves and get your hands dirty...

OpenWrt is already years ahead of what you can buy at Best Buy.

# In Disaster, There is Opportunity



Every line of code is available to modify; changes that work go upstream to OpenWrt and Linux as fast as are validated

CeroWrt is an advanced build of OpenWrt, using WNDR 3700v2 and WNDR3800 routers for more flash, Atheros radios, and fast CPU

Today running Linux 3.3.8 release with CoDel, BQL. Running fq\_codel on WiFi, which is today only partially effective due to buffering in the drivers due to 802.11n aggregation

Current Bind & DNSsec in chroot jail; dnsmasq also available

Routes, not bridges; 6 networks in the box

Real web server, proxy, IPv6 support, mesh networking, extensive network test tools, etc.....

Come help test, develop, and improve

Demonstrate your heretical ideas with running code!



# Remember, we are all in this bloat together!

Please come help before we sink!

My Blog - http://gettys.wordpress.com

Other Information

http://www.bufferbloat.net/projects/bloat

