Congestion Control Workshop

Constraints

July 28, 2012

Goals

- Goal of this session is to have a discussion about observations paper authors have made regarding the deployment status and constraints that impact the solutions.
- Want to increase understanding of topics like:
  - Too large buffers cause problems for real-time traffic
  - The entire community has been stuck on bandwidth for over a decade.
  - Real-time traffic suffers from queue sharing
  - Interworking between different layers is needed

Latency

150ms	**Historical "upper bound" on latency for VoIP**
100ms	**Echoing characters is hard to perceive as delayed**
100ms	RTT Jim's house to San Francisco via Comcast
75ms	RTT Boston/San Francisco via a good path
20ms	Remote rubber banding is "solid" and feels attached
16ms	60hz refresh rate latency
13ms	One full sized TCP packet @ 1Mbps
7.5ms	Jim's minimum cable RTT
< 5ms	Gamers (to reduce probability of getting "fragged")

Latency, cont.

- Conclusion: You can't give away ***any time*** unnecessarily. Period. You must ***always*** be minimising the latency, everywhere. Your latency budget is *always* overdrawn for many paths.

The Network Edge

- Packet trains arriving at line rate at the edge of the network have become very long. and head of line blocking causes major latency spikes.
- Breaking up these bursts has become necessary. Edge devices (broadband gear, home routers) must solve this mess.
- Causes: smart hardware, bandwidth uber-alles benchmarking, nasty applications, sharded web sites, ack merging in broadband gear destroying ack clocking, etc....

Buffers

- Buffers are dimensioned incorrectly and it is not possible to statically size buffers, both for inherent delay properties and because the bandwidth is continually varying.
- BitTorrent-style applications fill uplink queues without the user realizing it.
- The number of devices sharing Cable/DSL links keeps increasing.
- Mechanisms to diagnose buffering problems aren't currently adequate.
- Shared links means that other devices fill queues without user's realising it.

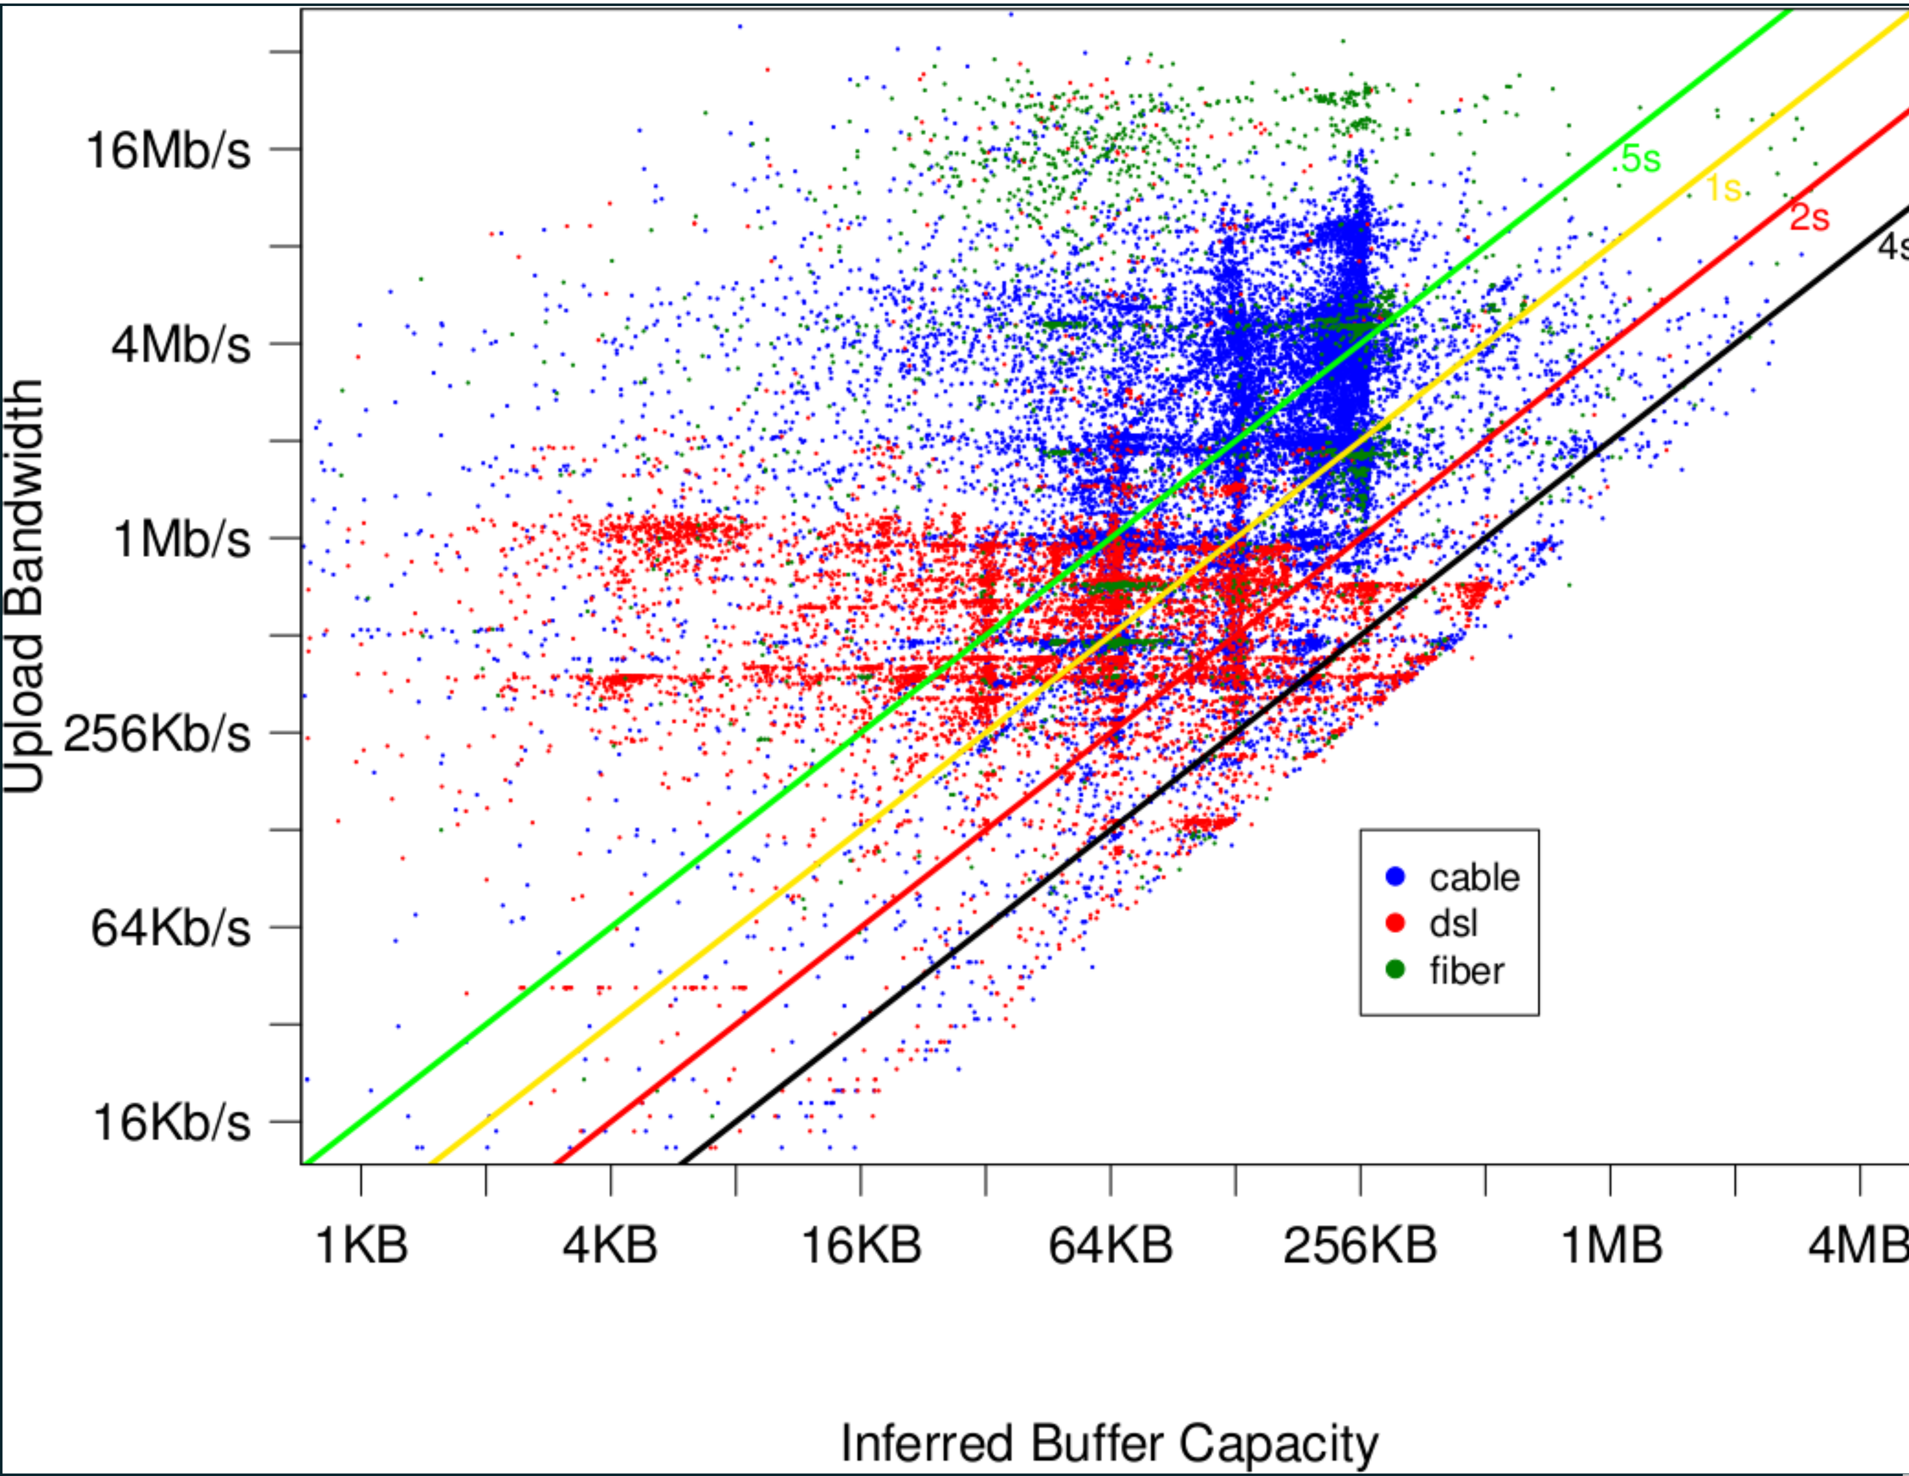

From: http: gettys.files.wordpress.com 2010 12 uplink_buffer_all.png

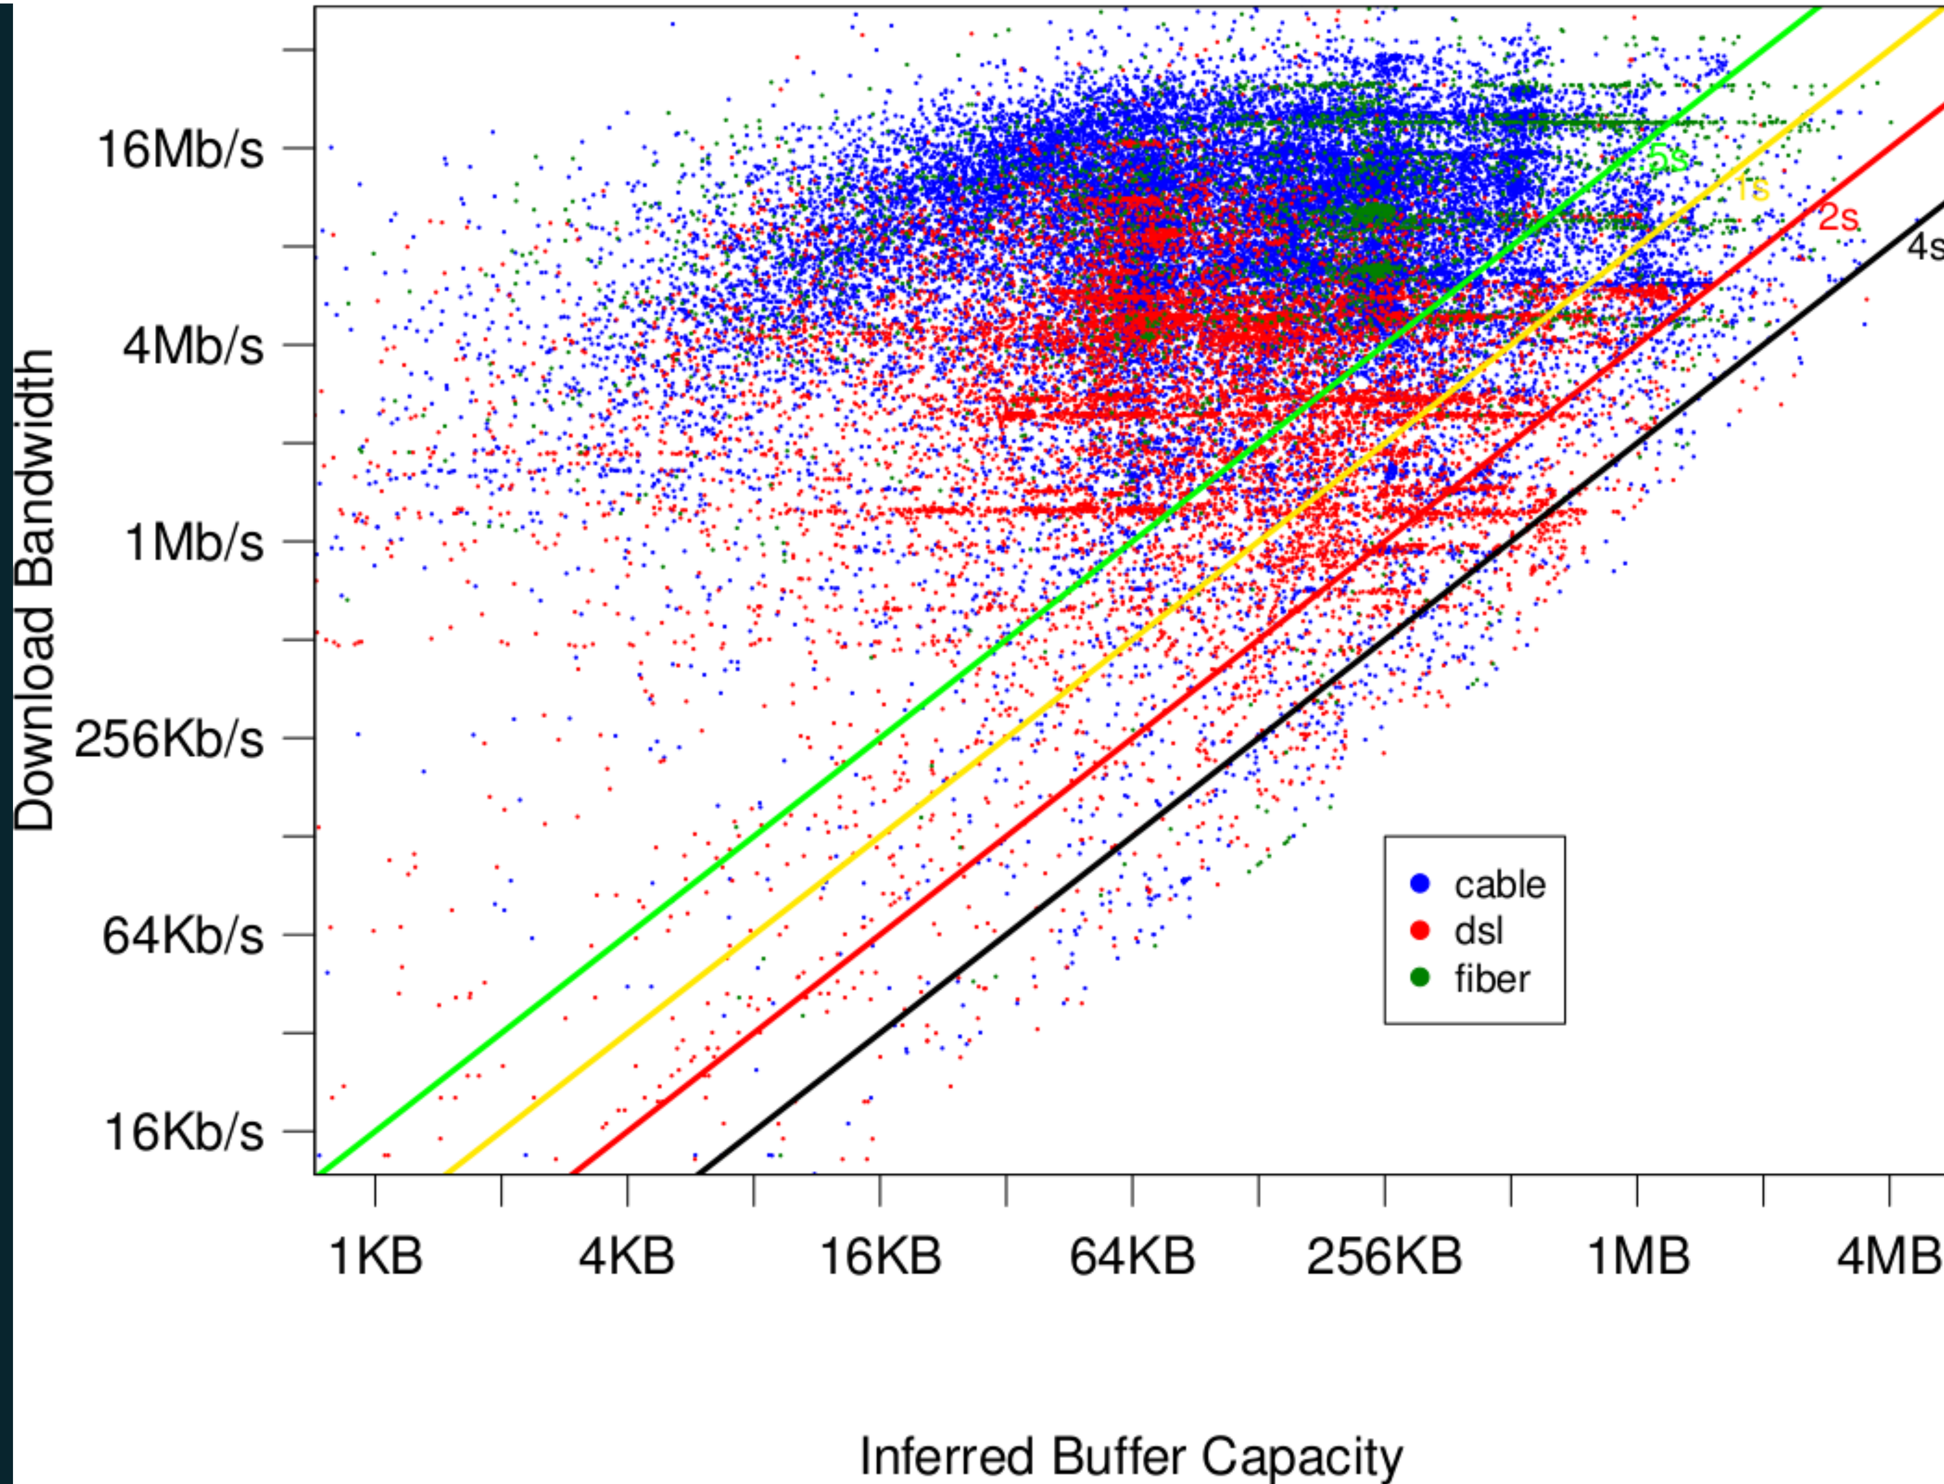

From: http: gettys.files.wordpress.com 2010 12 downlink_buffer_all.png

Recognizing Buffer Problems?

- Operators currently put stock in independent test results (such as speedtest.net) that measure TCP throughput on a long-running flow, and round-trip latency on an unloaded link.
- They do not routinely look at statistics of round-trip latency on a loaded link and sometimes dismiss results suggesting that they are capable of buffering > 5 seconds of data before dropping a packet.

Queue Sharing

- Competing TCP flows need a Round-Trip-Time of buffering.
- Without separating RT traffic from TCP in separate queues unacceptable delay or delay jitter is the result.
- It is the uncoordinated congestion management which causes the problems which stem from queue-sharing.
- Broadband service has but a single queue available for all of a customer's data traffic. Only ISP's telephony services today have access to other queues, despite the hardware being capable of multiple queues.
- End-to-End QoS is not deployable, but QoS at the ISP/sender and ISP/receiver interfaces might be enough.

Layer Interworking

- John: “We're acting as if we believe congestion will magically be solved by a new transport algorithm. It won't.”
  - We need action at the Network layer, the Transport layer, and the Application layer, at a minimum.
  - The application layer is the only practical place to balance what piece(s) to constrain to lower bandwidths.
- All flows relating to a user session should have a common congestion controller.

  For many applications, audio is much more critical than video. For example, video may back off, but the audio remains un changed.