Residential Broadband Networks: Characteristics and Implications

Marcel Dischinger, Andreas Haeberlen, Krishna P. Gummadi, and Stefan Saroiu†
Max-Planck-Institute for Software Systems, †University of Toronto

1. Motivation

- **Residential broadband networks are widely used**
  - >150 million users world-wide, >50% of Internet hosts in the U.S.
- **They are critical for emerging Internet applications**
  - Such as VoIP, music and movie downloads, and P2P systems
- **Yet, we understand very little about broadband networks**
  - They are very different from academic and research networks
- **There is little data characterizing the networks at scale**
  - Researchers have limited ability to measure commercial networks

2. Measurement Methodology

- **We measured 1500 hosts from 11 major cable/DSL ISPs**
  - 2 orders of magnitude larger than prior studies
- **We require no cooperation from the broadband hosts**
  - This allows our measurements to scale
- **Our methodology relies on three key observations:**
  - Many residential hosts respond to TCP/ICMP packet probes
  - We can infer several network characteristics using probe trains
  - We can find lots of residential IPs by crawling large P2P networks

3. Characteristics

- **Broadband network characteristics differ from the conventional wisdom about academic network properties**

A. Link Bandwidths

Downstream link bandwidths

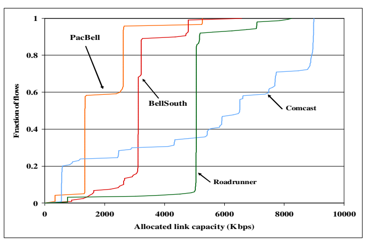

Upstream link bandwidths

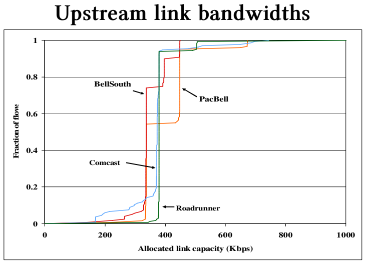

- ISPs allocate large downstream bandwidths, but tightly limit upstream bandwidths
- Measured link bandwidths closely match advertised rates for most ISPs. Some cable ISPs like Comcast do not allocate well-defined bandwidths

Bandwidth stability

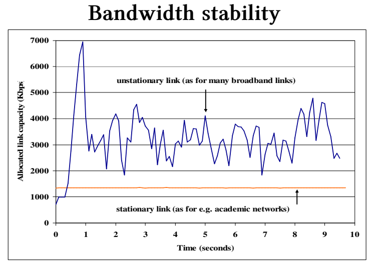

- Unlike DSL, cable link bandwidths vary considerably over time
- Likely causes: TDMA and link sharing

Traffic shaping / rate limiting

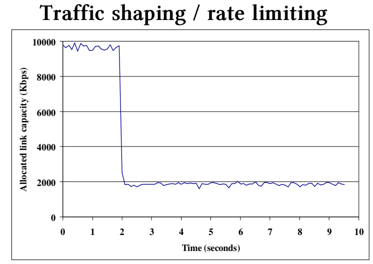

- Some ISPs allow short bursts of traffic to speed up small Web page downloads
- Other rate limit large flows

B. Packet Latencies and Losses

Latency variation

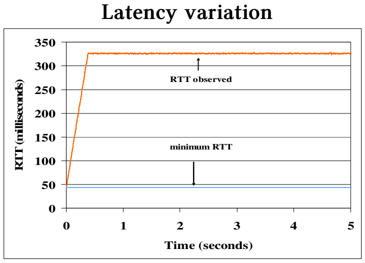

- Packet latencies can fluctuate wildly due to queuing delays
- RTTs can vary by a factor of 10 or more

Downstream queue sizes

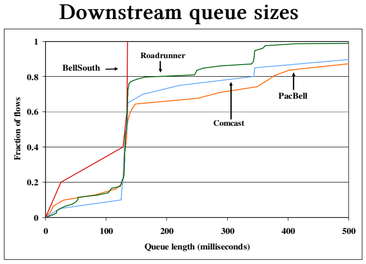

- Downstream queues can be significantly larger than one BDP
- Typical size 150 ms >> 50 ms typical RTT

Upstream queue sizes

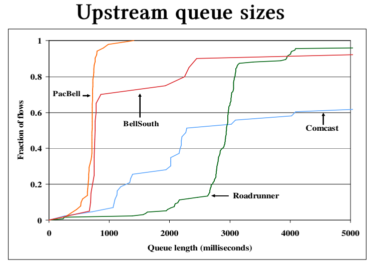

- Massive upstream queues, especially for cable links
- Typical sizes are in the order of seconds

Non-tail drop policies

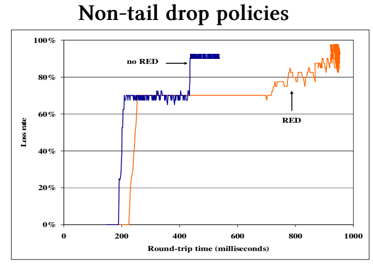

- Some ISPs use non-trail drop policies like random early detection (RED)

4. Implications

- **Our study has important implications for the design of future protocols and applications**
  - Increasing bandwidth asymmetry limits the effectiveness of P2P collaborative content distribution systems
  - Network coordinate systems like Vivaldi and GNP have to cope with highly varying path latencies
  - TCP is a bad gamble: Small window sizes limit throughput, large window sizes lead to large queuing delays. Hard choice between the throughput of bulk flows (BitTorrent) and the latency of interactive traffic (Skype)
- **Our trace-driven models of broadband links can help study their implications in simulations**

Skype can suffer from high delays