

Bandwidth Estimation in Broadband Access Networks

Venkat Padmanabhan
Microsoft Research

Joint work with:

Karthik Lakshminarayanan (Berkeley)

Jitu Padhye (MSR)


Bandwidth Estimation

- Several notions of bandwidth
 - bottleneck bandwidth: raw capacity
 - **available bandwidth**: spare capacity or “headroom”
 - other notions: fair share, bulk transfer capacity
- Desirable attributes of an estimation scheme
 - accurate, fast, lightweight, end-host-based


Previous Work

- Active area of research for 15+ years
- Focus primarily on wide-area wired paths
- Capacity Estimation
 - classical packet-pair
 - refinements: bprobe, nettimer, pathrate
- Available bandwidth estimation
 - early work: cprobe
 - packet rate method (PRM): PTR, pathload
 - packet gap method (PGM): IGI, spruce

Packet Rate Method: Pathload




Probing rate $<$ available bandwidth \Rightarrow no trend in OWD




Probing rate $>$ available bandwidth \Rightarrow increasing OWD

Packet Gap Method: Spruce



$d_i = d_o \Rightarrow$ no cross-traffic



$d_i < d_o \Rightarrow$ cross-traffic = $C^*(d_o - d_i)/d_i$

Traditional Link Model

- Assumptions made in previous work:
 - well-defined capacity
 - FIFO scheduling of probes and cross-traffic
 - fluid cross-traffic (infinitesimal packet size)
- But these assumptions break down in broadband access networks
 - cable modem, wireless (802.11), *DSL*

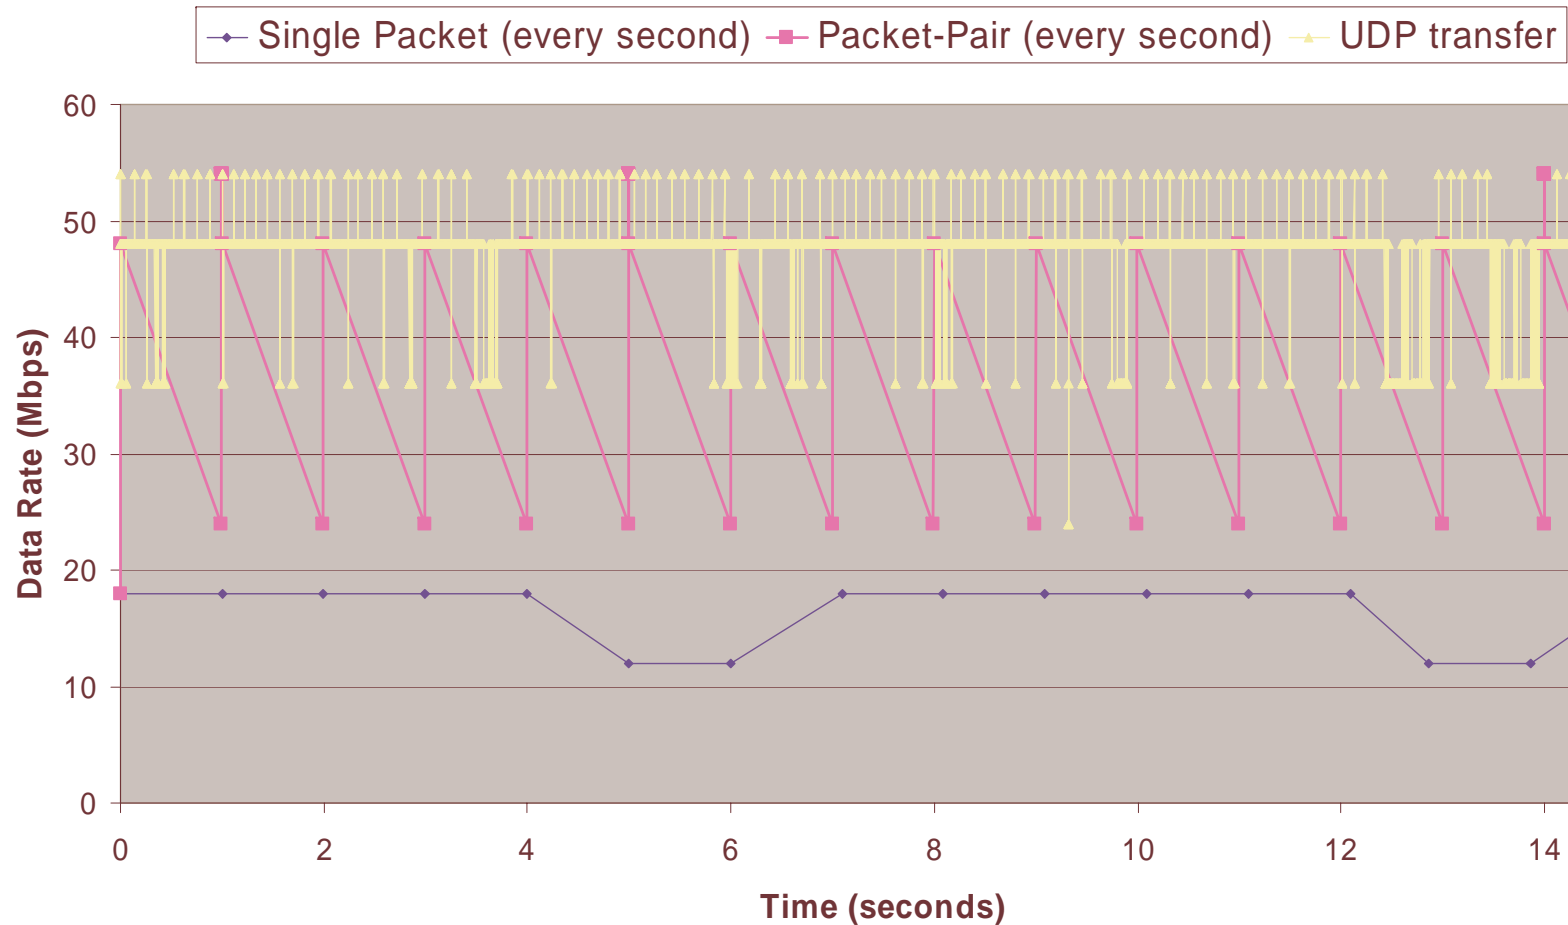
Why is broadband different?

- Lower cost \Rightarrow shared medium
- Pricing flexibility \Rightarrow “managed” links
- Variable channel quality \Rightarrow PHY adaptation

Specific Issues


- Link may not have well-defined capacity
 - rate regulation (e.g., token bucket)
 - dynamic multirate
- Non-FIFO scheduling
 - frame-level contention
- Interference between multirate links
 - bursty cross-traffic

#1: Dynamic Multirate



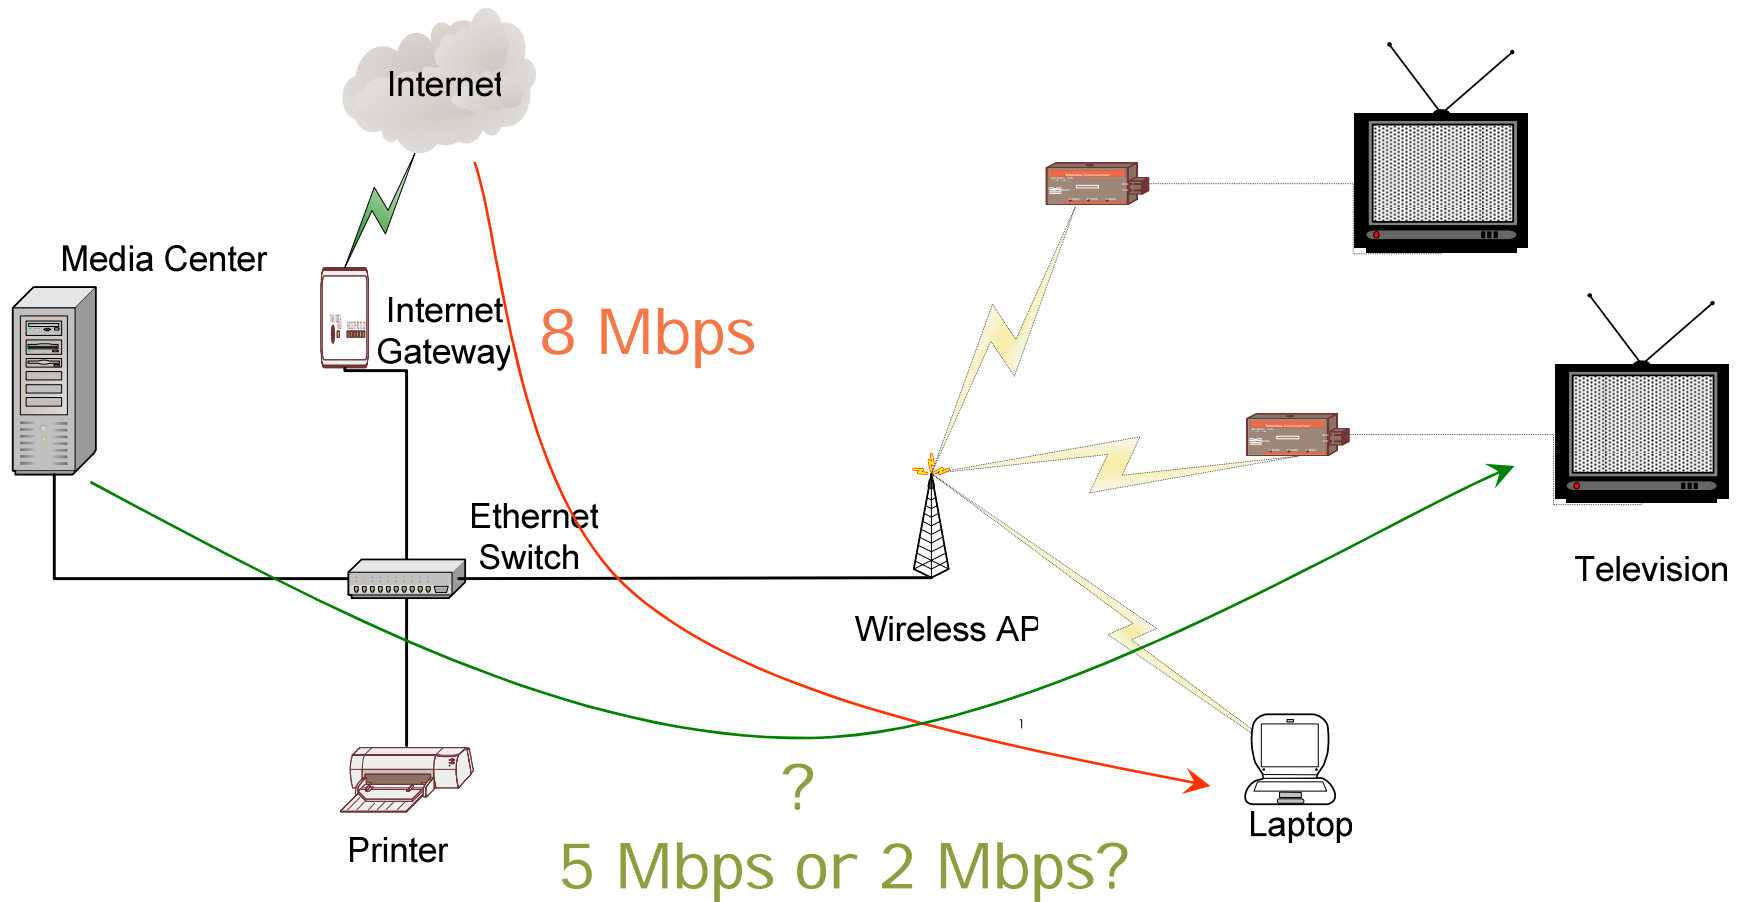
Rate seen by probes may not match that seen by traffic

#2: Non-FIFO Scheduling




Tendency towards estimating fair share on a frame-level basis

Is available bandwidth still interesting?



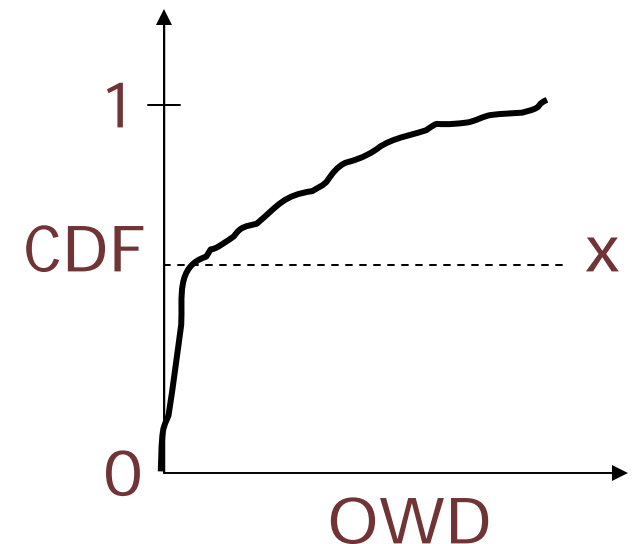
#3: Interference between Multirate Links



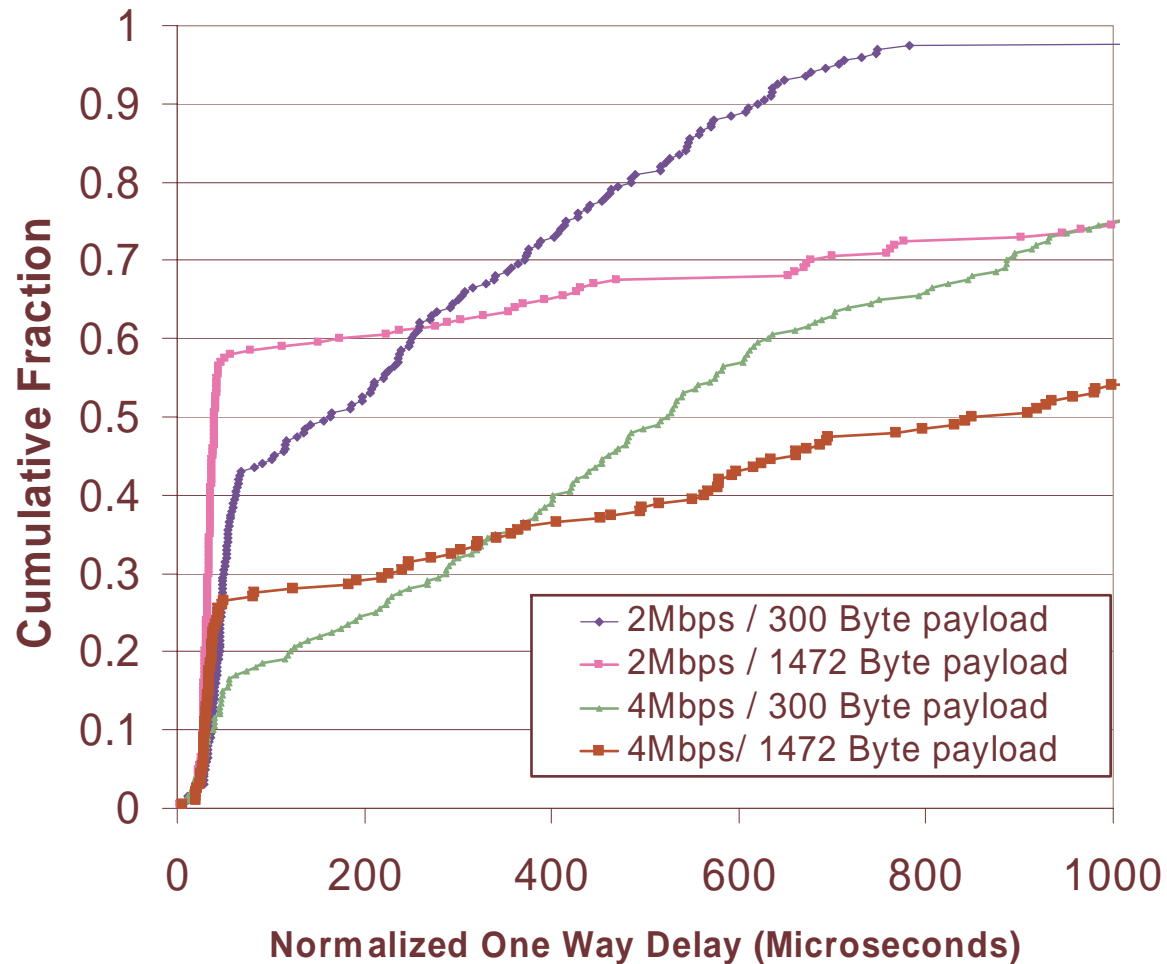
Burstiness of cross-traffic makes sampling hard

ProbeGap

- Goal: address non-FIFO scheduling & CT burstiness
- Key idea: probe for idle periods ("gaps") in the link
 - gather one-way delay (OWD) samples
 - "normalize" to eliminate clock offset
 - knee in CDF identifies idle fraction
 - multiply by capacity to estimate available bandwidth
- Issues
 - lightweight
 - 100-200 20-byte probes
 - clock skew
 - delays at other links

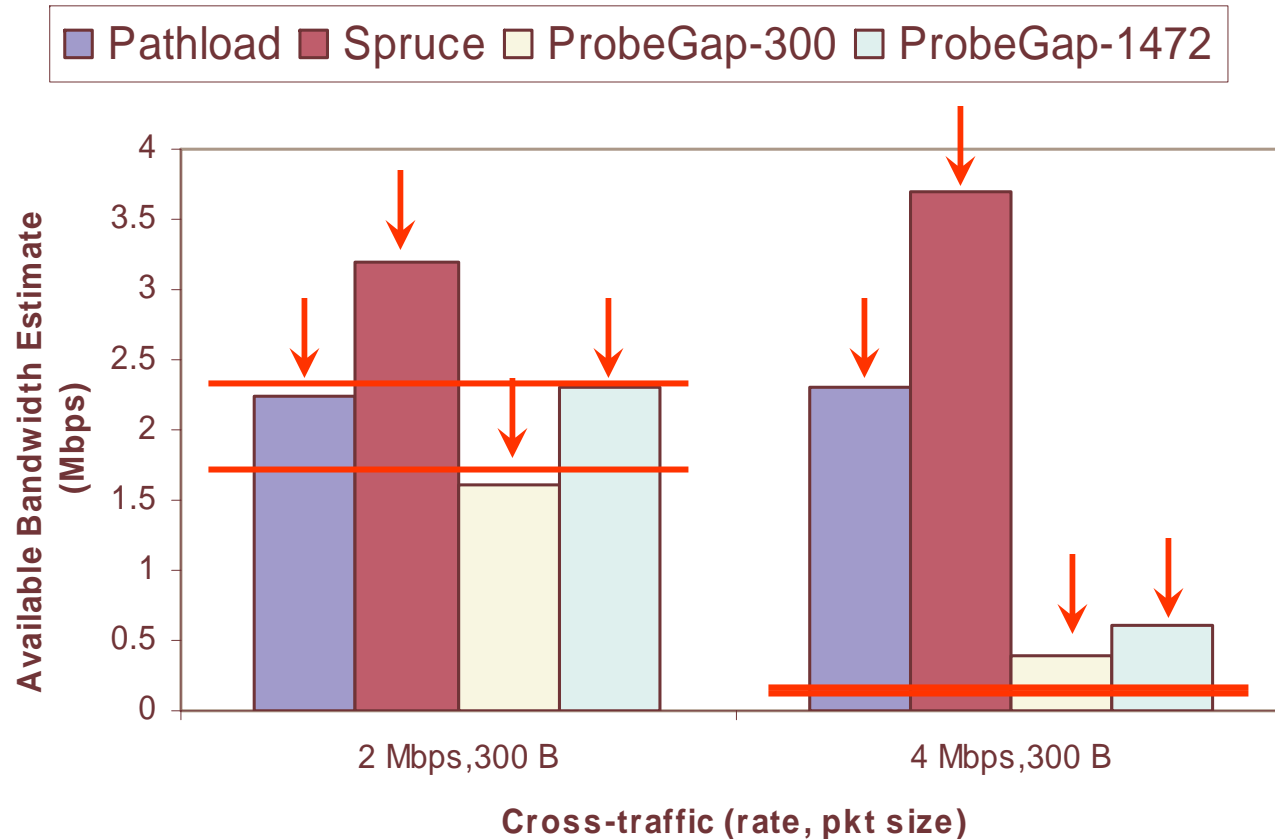


ProbeGap (802.11a, 6 Mbps)



Overestimation of idle fraction at high loads

802.11a, 6 Mbps single-rate



Overestimation due to tendency towards fair share (Pathload) and differential packet size (Spruce).
Probegap overestimates slightly under high load.

Conclusion

- Broadband access networks present new challenges to bandwidth estimation
 - performance experienced by probes may not be indicative of true performance
 - tendency to estimate fair share rather than available bandwidth
- ProbeGap looks promising in addressing some of the issues