



DASU - Characterizing Broadband Services

Zachary S. Bischof Mario A. Sánchez John S. Otto
John P. Rula Fabián E. Bustamante
Northwestern University



AquaLab

www.aqualab.cs.northwestern.edu/projects/dasu



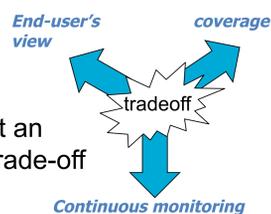
NORTHWESTERN UNIVERSITY

Benchmarking Broadband

ISPs have increased broadband coverage and performance in the past decade, but our ability to measure these networks has not kept pace

Several systems have been motivated by this observation – e.g. Netyalyzer, BISmark, SamKnows

Existing approaches exhibit an apparently unavoidable trade-off



Design Requirements

Broadband benchmarking should be...

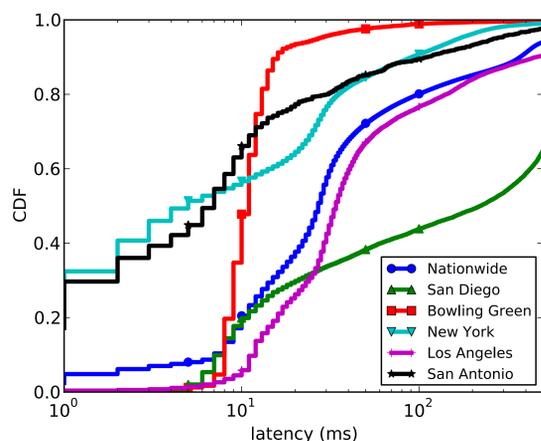
- (i) **at scale**, to capture the diversity of available providers and their services across locations
- (ii) **continuous**, to capture variations in service and dynamic changes in management policies
- (iii) **edge-based**, to capture a complete context, guarantee accuracy, and reduce bias

Our Approach – Dasu

- Network applications running on end systems meet these requirements
- Software-based model yields flexibility and low barrier to adoption
- Avoids the costs associated with hardware-based solutions
- By leveraging contextual information about the home network, Dasu attains the accuracy of hardware-based solutions

Scale

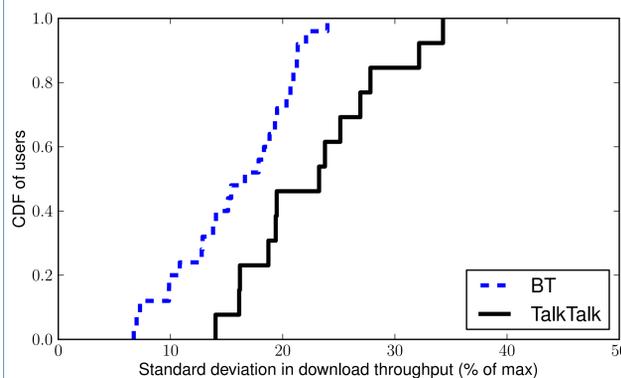
Disparate regions within an ISP have different characteristics (e.g. last mile latency distributions)



Diversity of last-mile latency measurements by city for customers of Time Warner Cable

Continuous

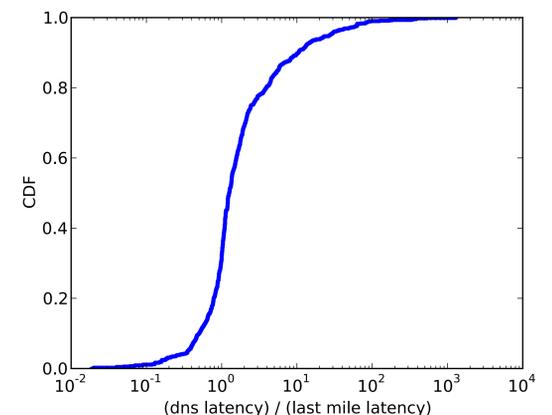
Observations over time are necessary to capture service variations and network outages



Variation in download throughput for users in BT's and TalkTalk's networks

Edge-based

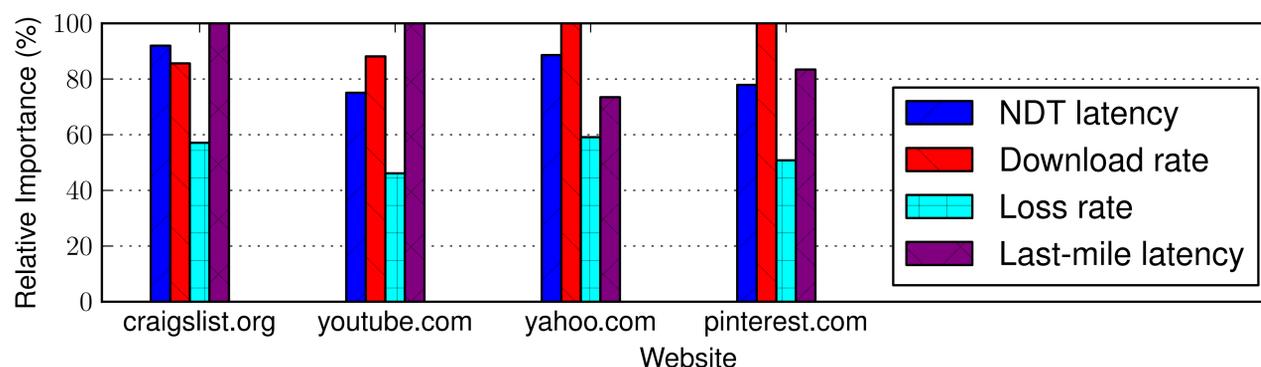
Many application-specific phenomena can only be seen from the network edge



Ratio of DNS latency to last mile latency (~40% of users have a DNS latency 2x higher than the last-mile)

Need for up the stack metrics

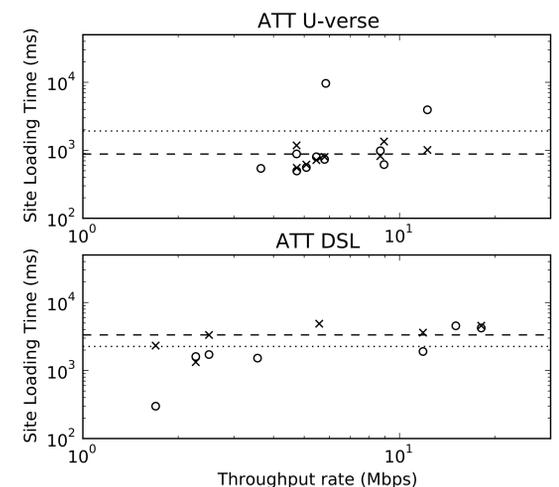
Adding measurements *up the stack*, from individual network-level metrics to objective quality of experience metrics, enables direct comparison of services



Relative importance of metrics when loading four popular websites (using gradient boosting)

No "silver-bullet" metric – craigslist.org is most influenced by latency while pinterest.com is most influenced by download throughput

Legend: craigslist.org mean (dotted), youtube.com mean (dashed), craigslist.org (circles), youtube.com (crosses)



For AT&T users, a faster DSL service does not translate into shorter page loading times, but switching access technologies does



Supported by NSF grants CNS 0644062, 0917233 & Google

