

latency for subscribers of AT&T's services with capacities above 10 Mbps across three geographic regions. While the difference between the South and Midwest distributions is relatively small (a KS-distance of 0.188), they both differ significantly from the West region – both have a KS distance above 0.40.

Overall, we observed a statistically significant difference in performance in 64% of comparisons when controlling for confounding factors. However, the average KS-distances of 0.29 was relatively low compared to the other factors in our analysis.

Subscription Rate Last, we look at the effect a user's subscription speed has on the connection's latency and packet loss performance. We found that 64% of comparisons between similar sets of users showed a statistically significant difference across capacity tiers. Interestingly, we found that as service capacity increased, the average network latency either decreased or stayed the same. We believe that this is likely due to the fact that download and upload capacity typically increased together and that the increased upload capacity minimizes the time waiting for transmission, resulting in a decreased latency. However, when looking at packet loss rates, only 28% of comparisons showed a significant difference.

5. RELATED WORK

In the following paragraphs we review the relevant literature looking at measurement location and bias in networks as well as previous attempts to measure and characterize residential broadband.

Active and passive monitor locations for network measurements. The location of active and passive monitors in networks has been studied extensively. One class of problem similar in nature to ours is the beacon placement problem within network tomography [5, 7]. Network tomography attempts to infer characteristics of the larger network from a collection of distributed monitors. Our goal is rather to characterize the overall performance of an end-host on a network edge, at the level of an ISP or autonomous system (AS) level.

Other avenues of research include the placement of passive monitors for recording and sampling network flows [3]. Suh et al. [9] studied the placement of these passive flow monitors within networks, discovering the NP-hardness of the problem, as well as the few total number of monitors needed to successfully record all flows within a network. The placement of network flow monitors are able to capitalize on common links traversed. Our work, in contrast, attempts to select measurement locations only from the network edge.

Residential Broadband Characterization. In addition to the governmental efforts by SamKnows and governmental agencies, several research efforts have looked into characterizing the quality and performance of residential broadband services. Sundaresan et al. [10] deployed instrumented home gateways to volunteer participants to perform network experiments through their BisMark system. Bischof et al. [2]

attempted crowdsourcing residential broadband characterization through the use of network intensive applications such as BitTorrent clients. These efforts, along with the ongoing government studies we analyzed in this paper, have provided a first look into the performance of residential broadband connections. Our goal is to leverage these initial research efforts to inform the next generation of research efforts.

6. CONCLUSION AND FUTURE WORK

This work examined the diversity within residential broadband networks, highlighting the importance of participant selection in existing, and future, deployments. We explored the underlying factors that cause heterogeneity in access networks, including variations across technologies and geographic regions within the same provider. We plan on expanding our analysis of network and vantage point diversity, leveraging additional broadband datasets and performing longitudinal analysis, with the goal of informing future vantage point selection through a principled approach.

As broadband services continue to grow in importance, attracting the attention of users and policy makers, there is a pressing need to determine the most appropriate metrics, measurement approaches and sampling strategies to help us derive a meaningful picture of their state. This work is a first but critical step toward that goal.

7. REFERENCES

- [1] Z. S. Bischof, F. E. Bustamante, and R. Stanojevic. Need, want, can afford – broadband markets and the behavior of users. In *Proc. IMC*, November 2014.
- [2] Z. S. Bischof, J. S. Otto, M. A. Sánchez, J. P. Rula, D. R. Choffnes, and F. E. Bustamante. Crowdsourcing ISP characterization to the network edge. In *Proc. of W-MUST*, 2011.
- [3] G. R. Cantieni, G. Iannaccone, C. Barakat, C. Diot, and P. Thiran. Reformulating the monitor placement problem: Optimal network-wide sampling. In *Proc. ACM CoNEXT*, 2006.
- [4] FCC. 2013 measuring broadband America February report. <http://data.fcc.gov/download/measuring-broadband-america/2013/Technical-Appendix-feb-2013.pdf>.
- [5] J. D. Horton and A. López-Ortiz. On the number of distributed measurement points for network tomography. In *Proc. IMC*, 2003.
- [6] S. S. Krishnan and R. K. Sitaraman. Video stream quality impacts viewer behavior: inferring causality using quasi-experimental designs. In *Proc. IMC*, 2012.
- [7] R. Kumar and J. Kaur. Efficient beacon placement for network tomography. In *Proc. IMC*, 2004.
- [8] SamKnows. Accurate broadband information for consumers, governments and ISPs. <http://www.samknows.com/>.
- [9] K. Suh, Y. Guo, J. Kurose, and D. Towsley. Locating network monitors: complexity, heuristics, and coverage. *Computer Communications*, 29(10), 2006.
- [10] S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè. Broadband internet performance: a view from the gateway. In *Proc. ACM SIGCOMM*, 2011.
- [11] S. Sundaresan, N. Feamster, R. Teixeira, and N. Magharei. Measuring and mitigating web performance bottlenecks in broadband access networks. In *Proc. IMC*, 2013.