

This experiment shows us that, even if we strive to put the WAN and the LAN connection on similar conditions, imposing the same RTT and PLR⁹, there is always a small bias in favor of the WAN connection. What happens is that, whenever there is a random fluctuation and a cycle without imposed losses ensues, connections run at their native speed, and the WAN connection, which is 10x faster, is able to deliver more bytes.

Indeed, as shown in Figure 6, which plots the outstanding window of the best run of each connection, the WAN outstanding window is free to grow much higher during slow start. Moreover, both connections have a three-second loss-free period, and, in this period, the WAN congestion window increases linearly, while the ADSL one seems to saturate, probably because the pipeline is already full, and the connection is just filling the home router buffer.

4. CONCLUSION

This paper introduces a simple model, based on the Mathis formula, to estimate, at the application level, the packet loss rate (PLR) experienced by a TCP connection. To perform the estimation, the application needs to (i) measure the goodput of a bulk TCP transfer that runs for a fixed amount of time (10 seconds) and (ii) estimate the round trip time (RTT), before or after the transfer. Testbed experiments show that the model predicts the imposed PLR value with a residual (small) dependence on the RTT. Controlled Internet experiments show that the model can reliably predict the PLR imposed to an ADSL connection. The model can detect the injection of losses on a WAN connection as well. However, since the estimation is less precise in this case, our research efforts are currently focused on refining the model for this scenario.

5. REFERENCES

- [1] cURL and libcurl. <http://curl.haxx.se/>.
- [2] Data from M-Lab Tools. <http://measurementlab.net/data>.
- [3] Glasnost: Test if your ISP is shaping your traffic. <http://broadband.mpi-sws.org/transparency/bttest-mlab.php>.
- [4] GNU Wget. <http://www.gnu.org/software/wget/>.
- [5] Measurement-Lab | M-Lab. <http://measurementlab.net/>.
- [6] netem | The Linux Foundation. <http://www.linuxfoundation.org/collaborate/workgroups/networking/netem>.
- [7] Network Diagnostic Tool (NDT). <http://www.internet2.edu/performance/ndt/>.
- [8] Neubot | The network neutrality bot. <http://neubot.org/>.
- [9] Neubot data | Neubot. <http://neubot.org/data>.
- [10] NPAD - Network Path and Application Diagnosis project. <http://www.ucar.edu/npad/>.
- [11] tcptrace - Official Homepage. <http://www.tcptrace.org/>.
- [12] Basso, S. and Servetti, A. and De Martin, J.C. Rationale, Design, and Implementation of the Network Neutrality Bot. In *Congresso AICA 2010 (L'Aquila)*, 2010.
- [13] Basso, S. and Servetti, A. and De Martin, J.C. The hitchhiker's guide to the Network Neutrality Bot test methodology. In *Congresso AICA 2011 (Torino)*, 2011.
- [14] Basso, S. and Servetti, A. and De Martin, J.C. The network neutrality bot architecture: a preliminary approach for self-monitoring of Internet access QoS. In *Computers and Communications (ISCC), 2011 IEEE Symposium on*, pages 1131–1136. IEEE, 2011.
- [15] Bauer, S. and Clark, D. and Lehr, W. Understanding broadband speed measurements. 2012.
- [16] Bischof, Z.S. and Otto, J.S. and Sánchez, M.A. and Rula, J.P. and Choffnes, D.R. and Bustamante, F.E. Crowdsourcing ISP characterization to the network edge. In *ACM SIGCOMM Workshop on Measurements Up the Stack (W-MUST)*. ACM, 2011.
- [17] Cardwell, N. and Savage, S. and Anderson, T. Modeling TCP latency. In *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 3, pages 1742–1751. IEEE, 2000.
- [18] De Martin, J.C. and Glorioso, A. The Neubot project: A collaborative approach to measuring internet neutrality. In *Technology and Society, 2008. ISTAS 2008. IEEE International Symposium on*, pages 1–4. IEEE, 2008.
- [19] Guillemin, F. and Robert, P. and Zwart, B. Performance of TCP in the presence of correlated packet loss. In *15th ITC Specialist Seminar on Internet Traffic Engineering and Traffic Management (Wurzburg)*, 2002.
- [20] Mathis, M. and Semke, J. and Mahdavi, J. and Ott, T. The macroscopic behavior of the TCP congestion avoidance algorithm. *ACM SIGCOMM Computer Communication Review*, 27(3):67–82, 1997.
- [21] Mellia, M. and Zhang, H. TCP model for short lived flows. *Communications Letters, IEEE*, 6(2):85–87, 2002.
- [22] Padhye, J. and Firoiu, V. and Towsley, D. and Kurose, J. Modeling TCP throughput: A simple model and its empirical validation. In *ACM SIGCOMM Computer Communication Review*, volume 28, pages 303–314. ACM, 1998.
- [23] Palfrey, J. and Zittrain, J. Better Data for a Better Internet. *Science*, 334(6060):1210–1211, 2011.

⁹These are similar conditions since the two connections experience the same RTT. And since, for both connections, the imposed PLR is much higher than the default one, and dominates over it.