Poster: The Path Less Traveled: Overcoming Tor's Bottlenecks with Traffic Splitting

Mashael AlSabah, PhD Student
Kevin Bauer, Post Doctoral Fellow
Tariq Elahi, PhD Student
Ian Goldberg, Associate Professor
Email: [malsabah,k4bauer,mtelahi,iang]@cs.uwaterloo.ca

Cheriton School of Computer Science
University of Waterloo
Waterloo, Ontario

I. Poster Abstract

Tor [6] is a widely used low-latency anonymity network, which offers strong privacy guarantees by tunnelling a user's Internet traffic through virtual circuits consisting of multiple intermediate overlay routers using a layered encryption scheme based on onion routing [11]. Beyond enabling anonymous communications online, Tor has become an essential tool in the fight against Internet censorship. Today, regimes around the world continue to aggressively filter [17], monitor [10], or explicitly block access [9] to certain types of online content. While Tor serves an estimated 400,000 users on a daily basis [15], its public infrastructure of over 3000 public relays can be easily blocked. In response, Tor uses special unlisted relays called *bridges* to aid users residing within regimes, such as China, that explicitly block the Tor network. Unfortunately, bridges generally provide a lower quality of service than Tor's public infrastructure.

Problem Statement Although Tor's primary goal is to support *real-time interactive* applications such as web browsing, the network suffers from a variety of performance problems [7] that are manifested as high and variable delays which result in a poor user experience. This poor experience discourages Tor's adoption and ultimately results in a smaller user base and weaker anonymity for all users [5].

Dynamic Traffic Splitting for Tor. In this work, we recognize that the diversity of bandwidth provided by Tor's volunteer-operated routers, and in particular the low-bandwidth bridges, degrades performance. We also recognize the significance of improving the performance of some high-throughput applications, such as streaming web videos, for Tor users. We propose an unconventional approach to improving performance when using low-bandwidth routers and bridges: *Tor users should split their traffic across multiple semi-disjoint circuits.*

In the context of Tor, traffic splitting offers the following benefits:

– *Improve load balancing.* When routers become over-utilized and experience congestion, splitting traffic across semi-disjoint paths can ease the burden on the congested circuit; under our scheme, circuits need only share a common exit router.

– *Improve performance with low-bandwidth relays.* By splitting data over multiple circuits, the user's throughput can achieve up to the aggregate throughput of all circuits rather than a single one. This is particularly useful when a circuit uses a low-bandwidth router. Tor's router selection algorithm favors routers that have higher bandwidths to ensure sufficient throughput to transport users' traffic and to balance the traffic load across Tor's routers. However, individual Tor routers can have vastly different bandwidth capacities, ranging from 20 KiB/s to over 20 MiB/s. Our results show a long-tailed distribution of download times for 1 MiB files over the course of two different months: January and October 2012. These slower downloads often correspond to circuits that used at least one low-bandwidth router. By combining multiple circuits with low-bandwidth nodes, the attainable throughput is no longer bound by the bottleneck node, but is instead the aggregate of each individual circuit's throughput.

Our approach. We design, implement, and evaluate *Conflux*,[1] a novel congestion-aware traffic splitting and load balancing algorithm for anonymous communication networks. Conflux forwards a client's individual cells down multiple circuits that share a common exit router as shown in Figure 1. Our algorithm dynamically measures the throughput of each constituent circuit and assigns traffic to each in proportion to its observed throughput. Our approach performs sub-stream traffic splitting, which provides a fine granularity of load balancing, as splitting can be performed at the individual cell level. This allows the traffic that is sent on a circuit to correspond to its desired load. The circuit's endpoints (the client and the exit router) are responsible for splitting the traffic at one endpoint and buffering, re-ordering, and delivering in-order cells to the application at the other end of the circuit. This approach can be deployed incrementally, as only clients and exit routers need

[1] Conflux: a flowing together of rivers or streams.

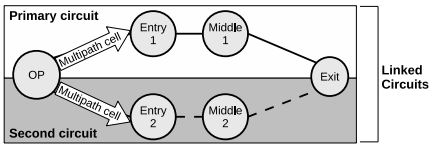

Fig. 1. Multipath construction and stream linking

to upgrade to support it.

To quantify the performance benefits of our proposed design, we perform a variety of live and whole-network experiments on an emulation-based Tor network testbed [2]. Our evaluation indicates that our approach can result in decreased queueing delays and increased throughput for users, particularly those who rely on low-bandwidth bridges to access the Tor network. We also find that, under light traffic loads, Conflux improves performance for clients who use Tor to access streaming videos (such as blocked YouTube videos[2]). Improving performance for such users is important, as streaming video websites are becoming a dominant source of Internet traffic [8], [12].

We also critically evaluate the security implications of utilizing additional circuits in light of the well-studied end-to-end traffic confirmation attack [13], [14]. Our analyses indicate that our scheme does not increase users' vulnerability to this attack, even when the adversary uses powerful selective denial of service tactics [1], [3], [16], [4] to maximize the number of circuits that can be compromised.

Contributions. This work offers these contributions.

1) To improve performance for bridge and streaming application users, we design, implement, and evaluate a dynamic traffic splitting scheme that distributes the traffic load across circuits according to each circuit's bandwidth capacity.
2) Our live performance analysis indicates that Conflux results in an expected improvement of 30% in a typical Tor client's queueing delay and up to 75% in total download time. Whole-network experiments show that noticeable improvements are possible even when most or all clients adopt Conflux.
3) We analyze the security of our scheme and argue that there are no additional risks to users' anonymity.

REFERENCES

[1] K. Bauer, D. McCoy, D. Grunwald, T. Kohno, and D. Sicker. Low-Resource Routing Attacks against Tor. In *Proceedings of the Workshop on Privacy in the Electronic Society (WPES 2007)*, pages 11–20, October 2007.

[2] K. Bauer, M. Sherr, D. McCoy, and D. Grunwald. ExperimenTor: A Testbed for Safe and Realistic Tor Experimentation. In *Proceedings of the USENIX Workshop on Cyber Security Experimentation and Test (CSET)*, pages 51–59, August 2011.

[3] N. Borisov, G. Danezis, P. Mittal, and P. Tabriz. Denial of Service or Denial of Security? How Attacks on Reliability can Compromise Anonymity. In *Proceedings of CCS 2007*, pages 92–102, October 2007.

[4] A. Das and N. Borisov. Securing Tor Tunnels under the Selective-DoS Attack. In *Proceedings of Financial Cryptography and Data Security*, February 2013.

[2]Note that while Tor's browser bundle disables Flash by default, it is now possible to stream videos over Tor using HTML5. We expect this use case of streaming video over Tor to increase in popularity in the near term.

[5] R. Dingledine and N. Mathewson. Anonymity Loves Company: Usability and the Network Effect. In *Workshop on the Economics of Information Security*, pages 547–559, June 2006.

[6] R. Dingledine, N. Mathewson, and P. Syverson. Tor: The Second-Generation Onion Router. In *Proceedings of the 13th USENIX Security Symposium*, pages 303–320. USENIX Association, 2004.

[7] R. Dingledine and S. Murdoch. Performance Improvements on Tor or, Why Tor is Slow and What We're Going to Do about It. http://www.torproject.org/press/presskit/2009-03-11-performance.pdf, March 2009.

[8] G. Maier, A. Feldmann, V. Paxson, and M. Allman. On Dominant Characteristics of Residential Broadband Internet Traffic. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*, pages 90–102, November 2009.

[9] The OpenNet Initiative. YouTube Censored: A Recent History. http://opennet.net/youtube-censored-a-recent-history. Accessed February 6, 2012.

[10] M. Piatek, T. Kohno, and A. Krishnamurthy. Challenges and Directions for Monitoring P2P File Sharing Networks-or: Why My Printer Received a DMCA Takedown Notice. In *Proceedings of the 3rd conference on Hot topics in security*, pages 12:1–12:7, July 2008.

[11] M. G. Reed, P. F. Syverson, and D. M. Goldschlag. Anonymous Connections and Onion Routing. *Selected Areas in Communications, IEEE Journal on*, 16(4):482 –494, May 1998.

[12] Sandvine. Sandvine Global Internet Phenomena Report — Fall 2011. http://www.sandvine.com/downloads/documents/10-26-2011_phenomena/Sandvine%20Global%20Internet%20Phenomena%20Report%20-%20Fall%202011.pdf, October 2011.

[13] A. Serjantov and P. Sewell. Passive Attack Analysis for Connection-Based Anonymity Systems. In *Proceedings of ESORICS*, pages 116–131, October 2003.

[14] V. Shmatikov and M.-H. Wang. Timing Analysis in Low-Latency Mix Networks: Attacks and Defenses. In *Proceedings of ESORICS 2006*, pages 18–33, September 2006.

[15] The Tor Project. Tor Metrics Portal: Users. http://metrics.torproject.org/users.html. Accessed November 2012.

[16] A. Tran, N. Hopper, and Y. Kim. Hashing It out in Public: Common Failure Modes of DHT-based Anonymitys Schemes. In *ACM Workshop on Privacy in the Electronic Society*, pages 71–80, November 2009.

[17] X. Xu, Z. M. Mao, and J. A. Halderman. Internet Censorship in China: Where Does the Filtering Occur? In *PAM*, pages 133–142, 2011.