Privacy-preserving Network Congestion Control

Hussein Darir Hussein Sibai Chester Cheng Sayan Mitra Geir Dullerud Nikita Borisov
University of Illinois at Urbana-Champaign

Project goals

A. *Develop algorithms and analysis tools for building congestion-aware traffic routing algorithms with provable privacy guarantees;*

B. *Develop the foundations, algorithms, and experimental systems for studying the trade-off between privacy and efficiency in different networks; and*

C. *Of particular interest are communication networks and other networks used for collection and dissemination of behavioral information.*

The Tor network

Tor Browser

- Our first step has been to study the problem of load-balancing in path selection in anonymous networks such as **Tor**.
- Users are increasingly turning to anonymous communication networks to protect themselves from surveillance and online tracking.

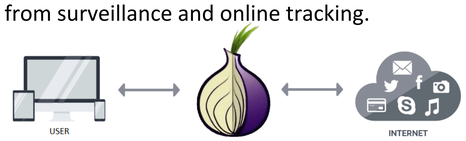

What is Tor?

- "Tor is free software and an open network that helps the user defend against traffic analysis, a form of network surveillance that threatens personal freedom and privacy."[1]

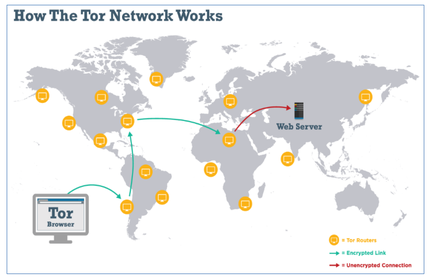

- To achieve anonymity in Tor, users' traffic is routed across a series of servers, called relays.
- Each user's path through the network, called a circuit, typically transits three of them.

Tor can be SLOW!

- Users choosing the paths imperfectly is a main reason.
- Currently relays are chosen randomly weighted by their estimated capacities.
- Current method of estimating capacity of relays is not accurate.

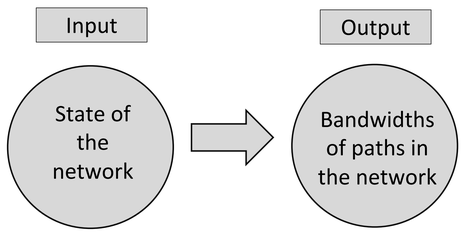

Figure. Max-min bandwidth allocation algorithm

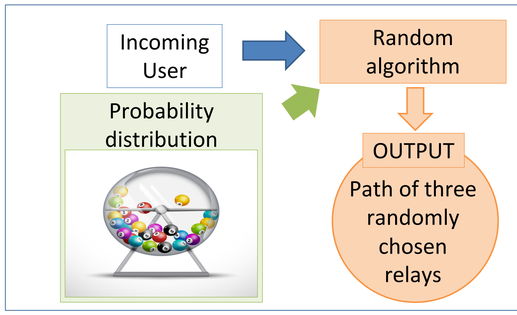

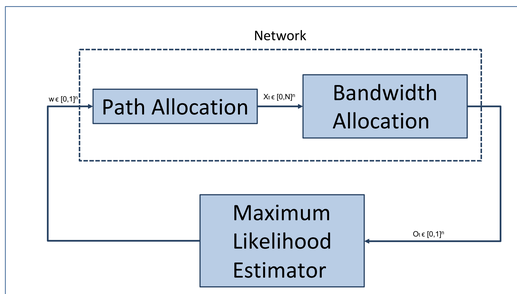

Relays capacities estimation

- Currently, a server periodically creates test paths that pass through all relays in the network and measures their allocated bandwidths.
- These bandwidths are then assumed to be the capacities of the corresponding relays that are released to the public.

This method can result in inaccurate measurements of the relays capacities
↓
Non-optimal allocation of loads on the relays
↓
Decrease the quality of service

- There were failed attempts to solve the problem using PID controller.

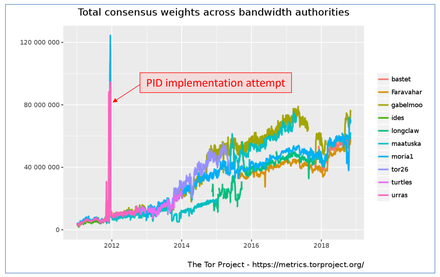

Shadow simulation

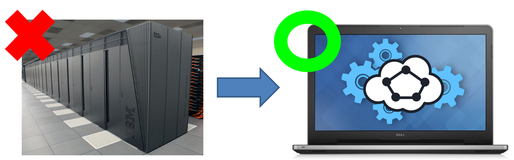

- Shadow creates an environment that allows simulated network connection between virtual nodes (clients, relays, and servers).
- Shadow runs Tor directly out of the box.
- Efficient network simulation in a single box.

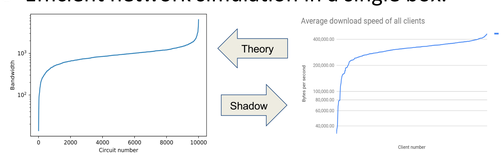

Contribution

- We are approaching the problem from a control theoretic and optimization point of view to release estimated capacities of the relays.
- We developed algorithms that result in provably accurate estimates of the relays capacities using maximum likelihood analysis.

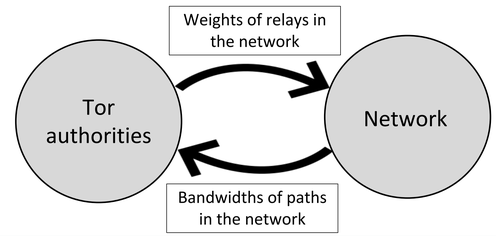

Simulations and results

- To make the problem tractable, we assume that each user path in the network consists of a single relay instead of three.
- Moreover, we ask the Tor authority to periodically allocate test paths with single relays to every relay in the network and measure the corresponding bandwidths.
- Given the measurements at a single period, the current capacities estimates and using maximum likelihood analysis, we derived a closed form solution for the optimal update of the estimates.

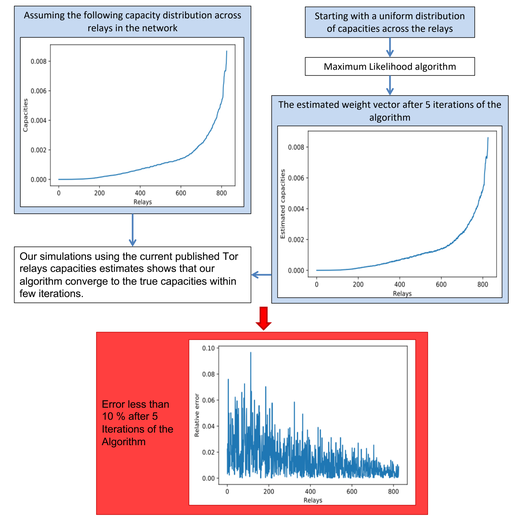

Project Website

https://wiki.illinois.edu/wiki/display/MitraResearch/Privacy-preserving+Network+Congestion+Control%3A+Theory+and+Applications

References

[1] https://www.torproject.org/
[2] Hussein Darir, Hussein Sibai, Nikita Borisov, Geir E. Dullerud, Sayan Mitra: TightRope: Towards Optimal Load-balancing of Paths in Anonymous Networks. In WPES '18: 2018 Workshop on Privacy in the Electronic Society, Oct. 15, 2018, Toronto, ON, Canada.
[3] https://metrics.torproject.org/

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. 1739966.