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.

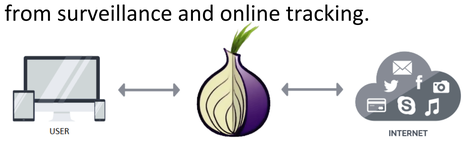

What is Tor?

- "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]

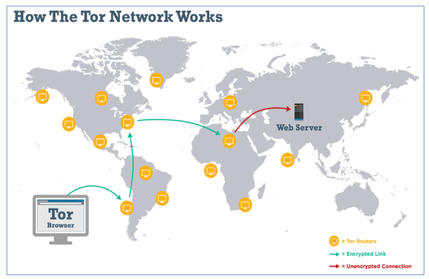

- To achieve anonymity in Tor, users' traffic is routed across a series of servers, called relays.
- Each user's path through the network, called a circuit, typically transits three of them.

Tor can be SLOW!

- Users choosing the paths imperfectly is a main reason.
- Currently relays are chosen randomly weighted by their estimated capacities.
- We demonstrate that this can create significant imbalances at any given point in time using flow-level simulations.

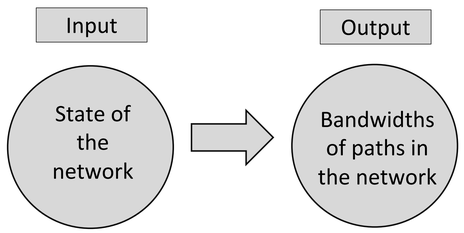

Figure. Max-min bandwidth allocation algorithm

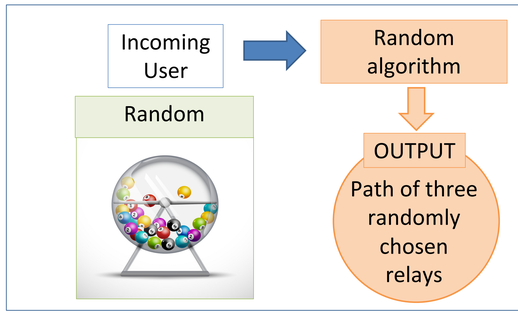

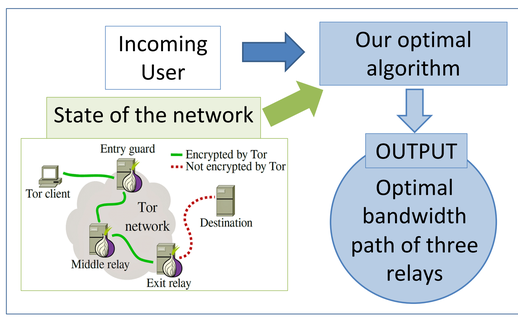

Contribution

- We modified the max-min allocation algorithm to select the three relays that would provide the maximum bandwidth for a new incoming user to the network.

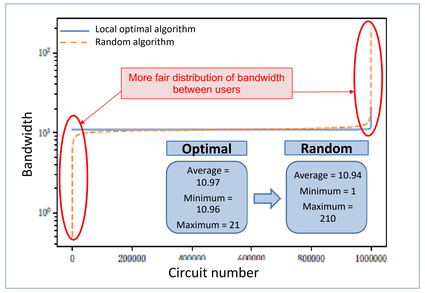

- We observe that the random algorithm for path allocation results in users that have very low bandwidth with respect to the average. While our locally optimal algorithm does not result in such cases.
- However, knowing the state of the network violates privacy.
- Instead of releasing the actual state for an incoming user to choose her path, a differentially private histogram representing the state of the network is released periodically for the public.
- For more details consult our paper [2].

Shadow simulation

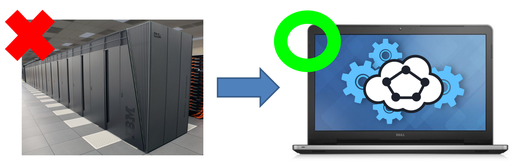

- Shadow creates an environment that allows simulated network connection between virtual nodes (clients, relays, and servers).
- Shadow runs Tor directly out of the box.
- Efficient network simulation in a single box.

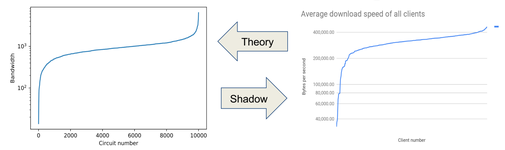

Relays capacities estimation

- Current method of estimating capacity of relays is not accurate.
- There were failed attempts to solve the problem using PID controller.

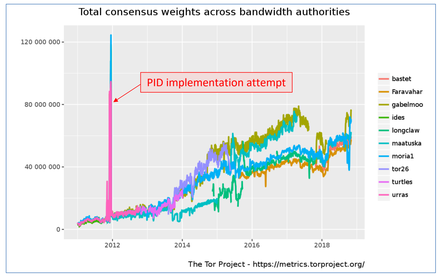

Ongoing work

Estimating relays capacities in Tor:

- In the previous project, the capacities of the relays were assumed to be known. However, this is not a realistic assumption.
- 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.
- However, this method can result in inaccurate measurements of the relays capacities.
- We are approaching the problem from a control theoretic and optimization point of view to release estimated capacities of the relays that ensures bandwidth fairness among users when selecting their paths using the random algorithm.

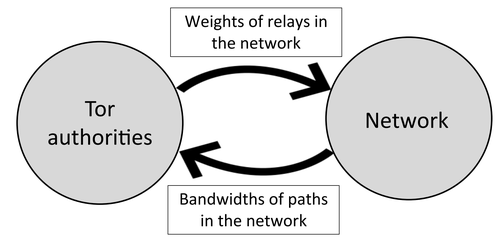

- Tor authorities wants to choose the weights such that they maximize the minimum bandwidth allocated to any path.

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.