Stony Brook University

האוניברסיטה העברית בירושלים
The Hebrew University of Jerusalem

Measuring and Mitigating AS-level Adversaries Against Tor

Oleksii Starov

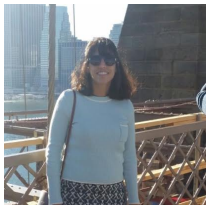

Adva Zair

Phillipa Gill

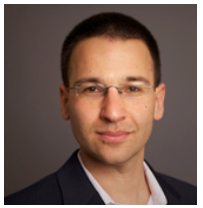

Michael Schapira

Rishab Nithyanand

Network-level Traffic Correlation Attacks

Internet routing is asymmetric.
Source -> Entry != Entry -> Source

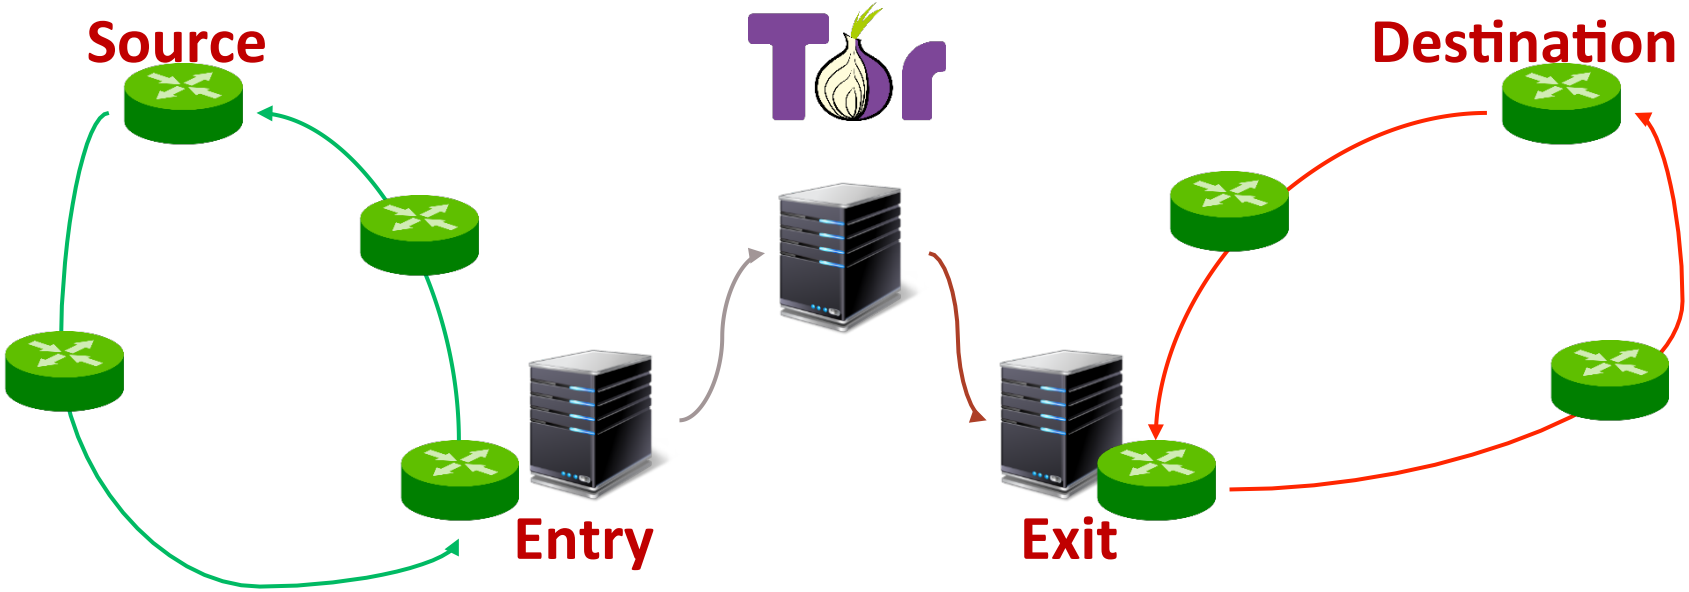

RAPTOR (USENIX Security 2015): Any AS on
(Source → Entry OR Entry → Source) AND (Exit → Dest OR Dest → Exit)
is in a position to launch a traffic correlation attack

Measuring Network-level Adversaries

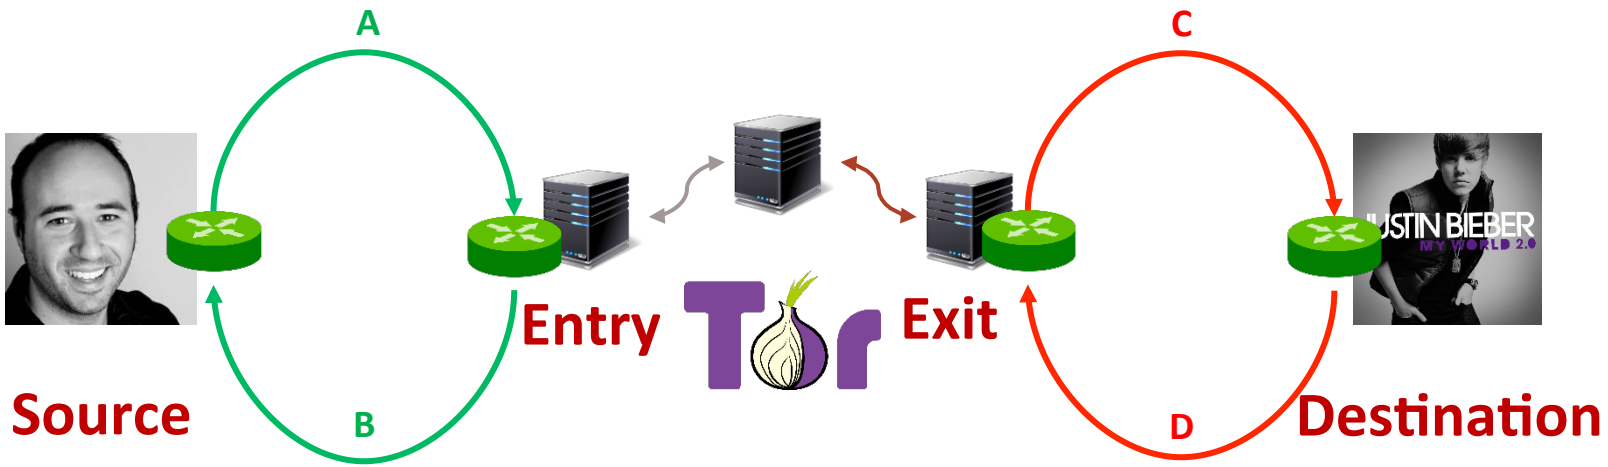

Goal: Quantify the threat from network-level adversaries

Approach: Identify ASes on A, B, C, and D

- ADV $= \{(A \cup B) \cap (C \cup D)\}$

Challenge: Traceroutes only let us obtain A

Measuring Network-level Adversaries

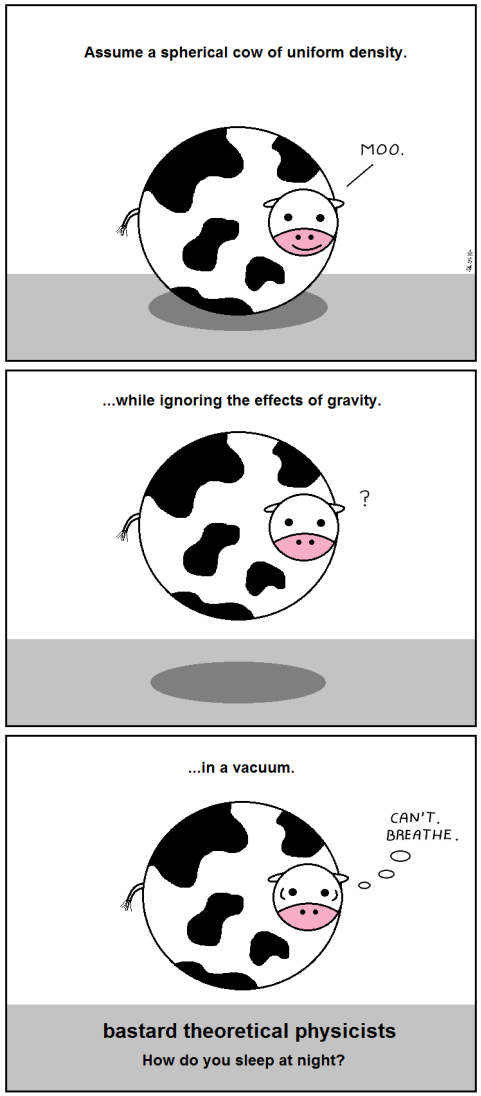

Our Approach: Spherical cows!

- Make assumptions about Internet routing.
- Obtain approximate AS-level paths.

Approximating ASes on a path (offline):

- AS Topology: **36K** ASes + **126K** relationships
- Use inter-AS relationships (customer, peer, provider) to decide whether an AS will route via another
  - Routing through customers > peers > providers, then prefer shortest paths
  - If there are multiple options, we consider all of them
- (see paper for validation)

Measuring Network-level Adversaries

10 Countries: BR, CN, DE, ES, FR, GB, IR, IT, RU, US

200 websites/country: Local Alexa T-100 + 100 Citizen Lab sensitive pages

Adversaries: Network-level, colluding network-level (see paper), and state-level

Measuring Network-level Adversaries

How vulnerable is vanilla Tor?

Main Circuit: Circuit carrying first “GET” request is vulnerable

Any Circuit: Circuit carrying any request is vulnerable

Network-level Adversary

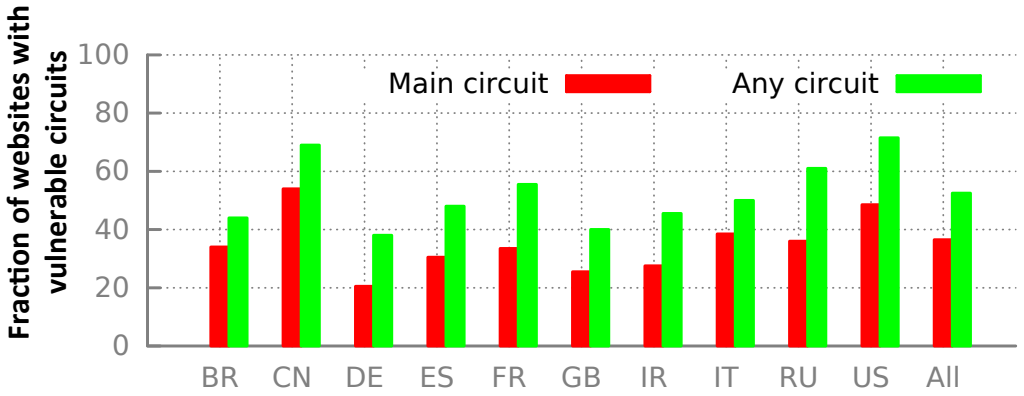

State-level Adversary

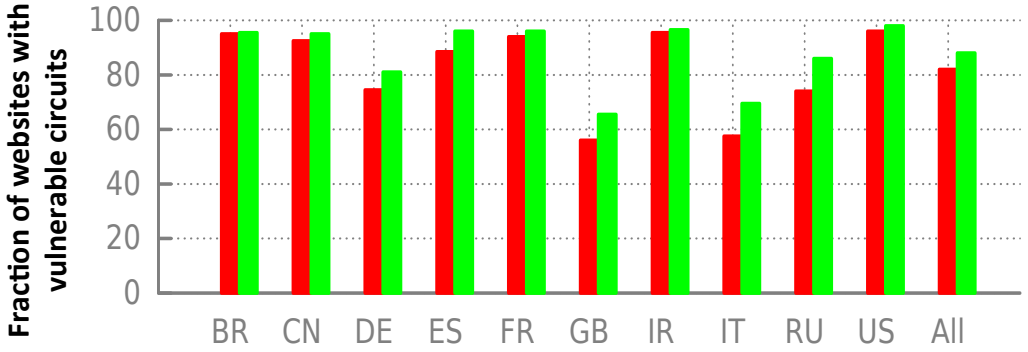

Measuring Network-level Adversaries

Can AS-aware relay selection help?

YES!

- > 20000 (source, destination) AS pairs in each country
- Consider 1000 * 1000 available (entry, exit) pairs
- **What fraction of the 20000 (source, destination) pairs have at most x% of their 1 million (entry, exit) pairs safe from network-level threats?**

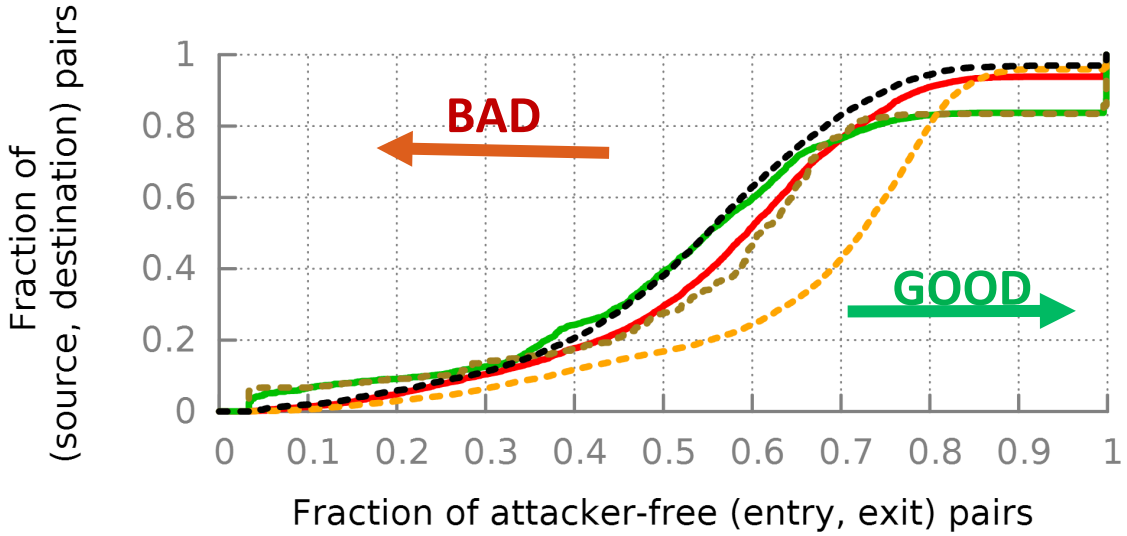

Astoria: This **AS**-aware **Tor** client **is a**lright

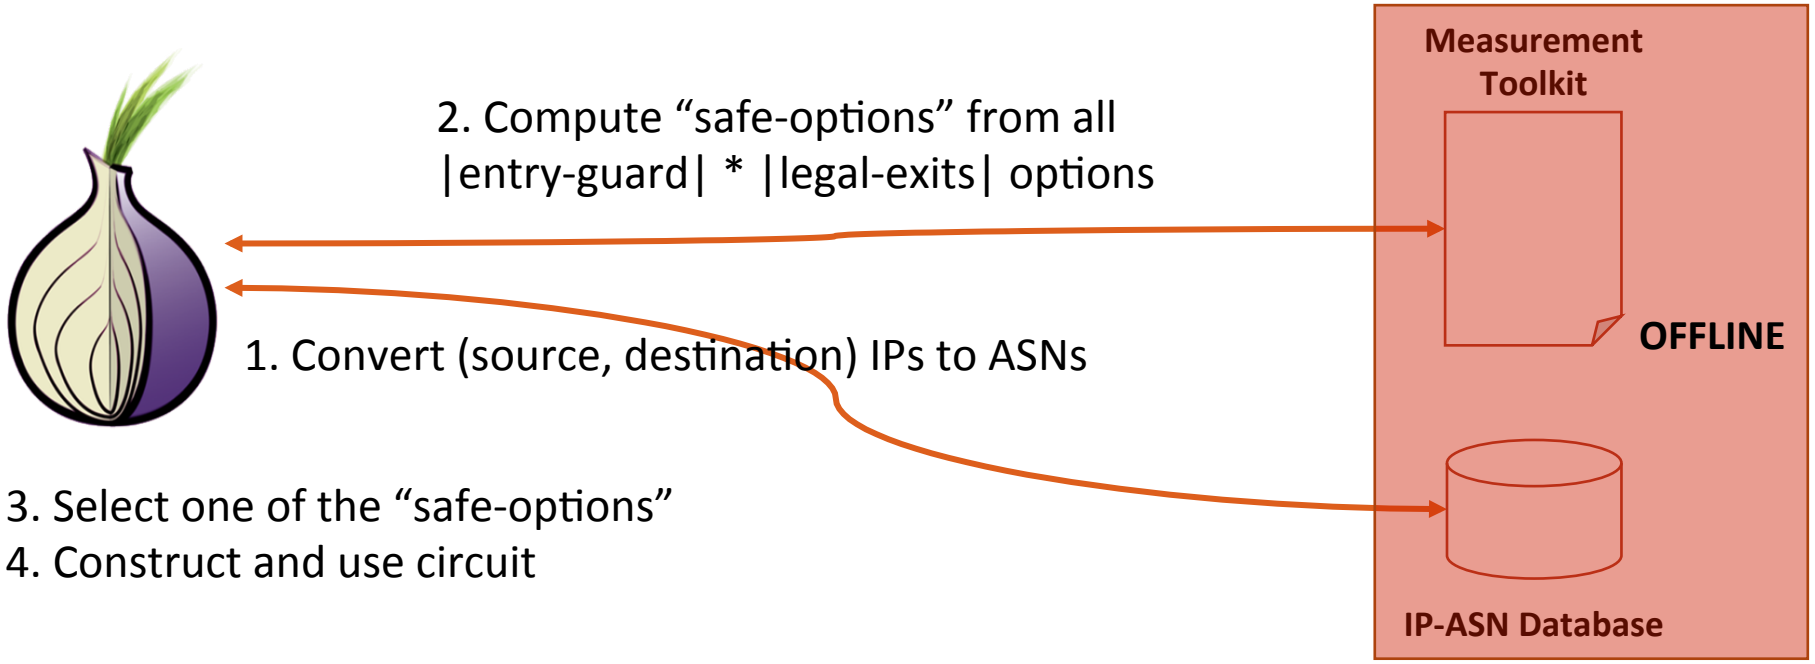

What if there are no safe options?

Astoria uses an LP to minimize number of circuits that are vulnerable to any single adversary. (see paper)

Astoria: Security Evaluation

Network-level Adversary

any: 53% -> 8%

main: 37% -> 3%

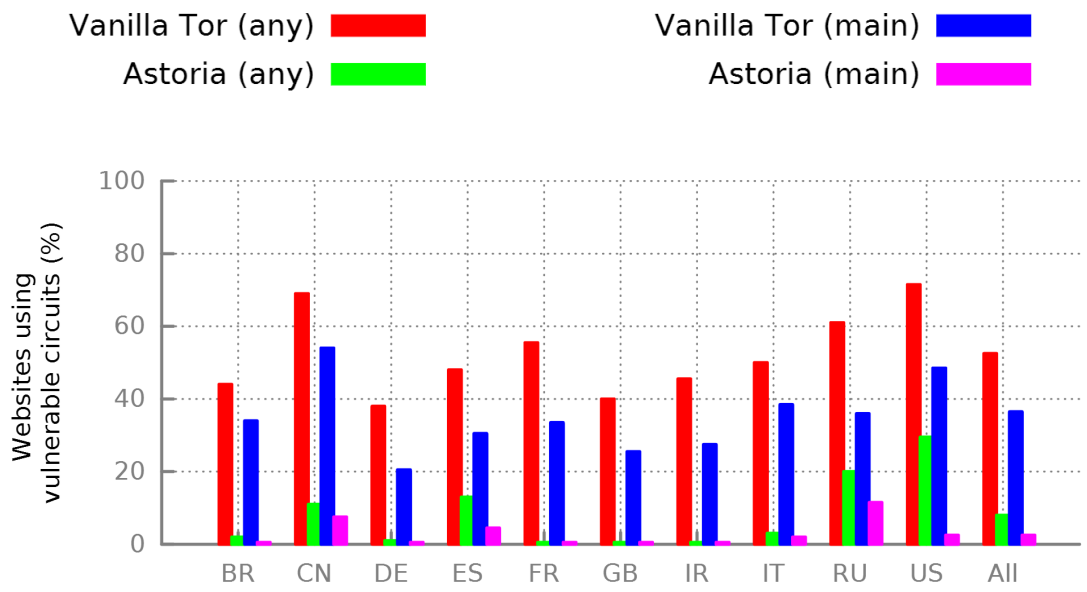

State-level Adversary

any: 88% -> 34%

main: 82% -> 27%

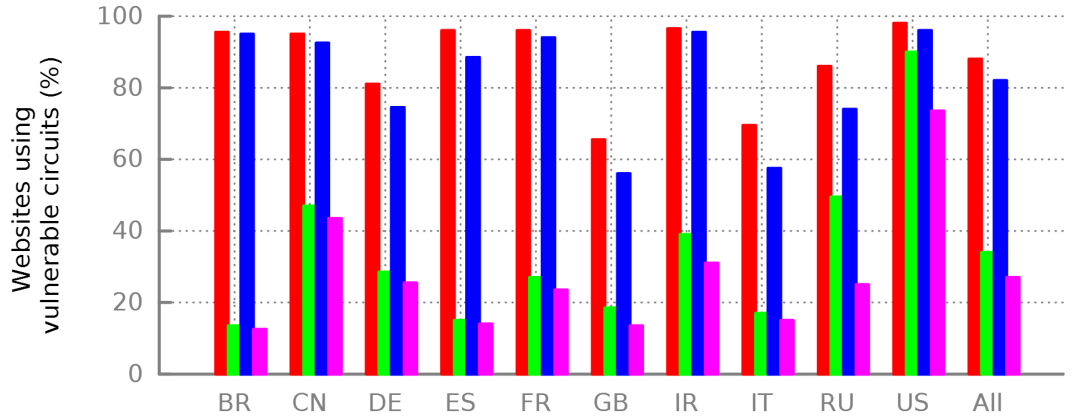

Astoria: Performance Evaluation

Page-load times

Tor: 5.9 sec

Astoria: 8.3 sec

Uniform: 15.6 sec

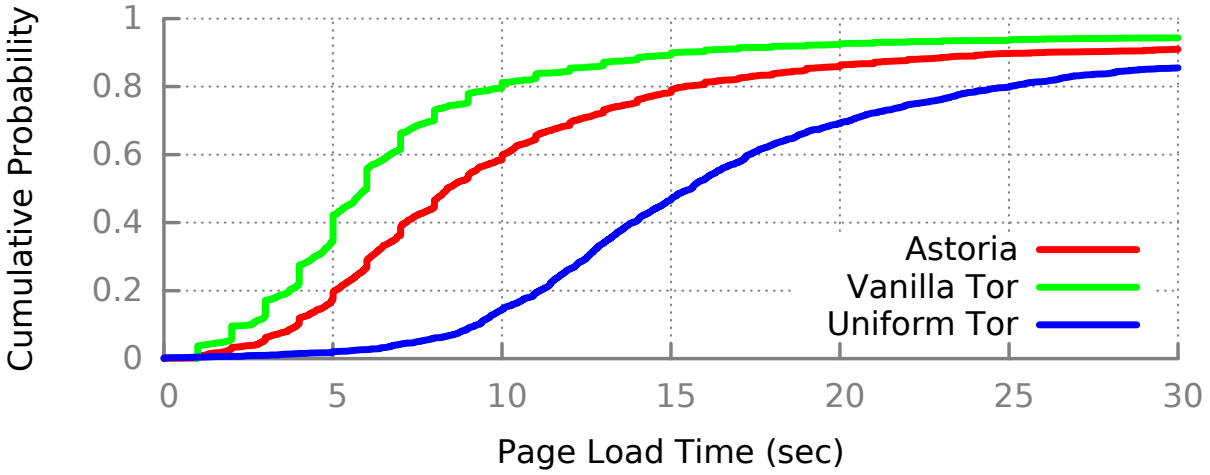

Load balancing

Similar to Tor*

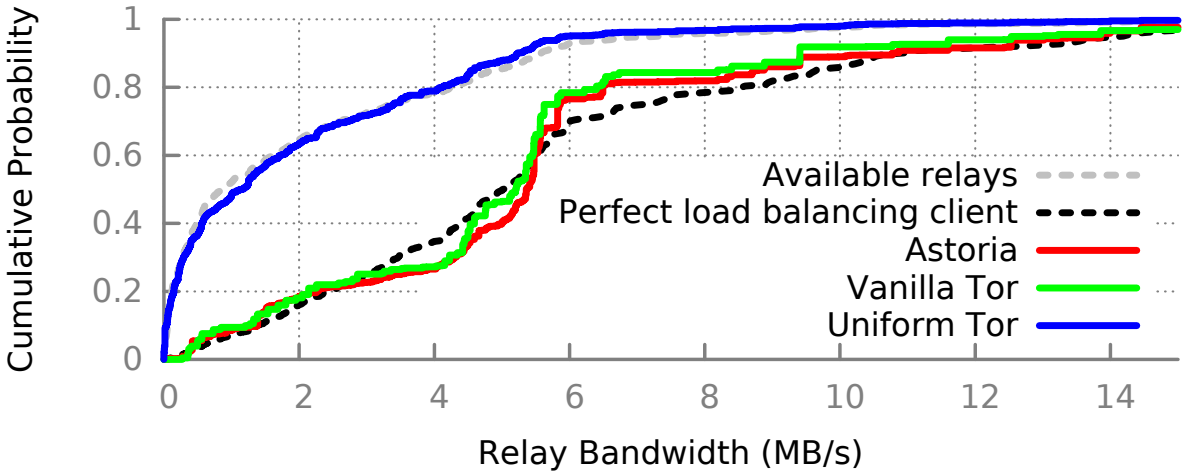

Conclusions

- Offline path-prediction toolkit to measure Tor vulnerability
- Significantly better security against network-level adversaries
  - Cuts number of vulnerable websites to less than 1/4th
  - Effectively deals with worst-case situations
- Load balancing: Similar to Tor
- Page-load times: Better than uniform, worse than Tor
  - Main problem: Cannot pre-build circuits like Tor
- Arguably weaker against relay-level adversaries (see paper)