



# PeerFlow: Secure Load Balancing in Tor

Aaron Johnson<sup>1</sup> Rob Jansen<sup>1</sup> Aaron Segal<sup>2</sup>

Nicholas Hopper<sup>3</sup> Paul Syverson<sup>1</sup>

<sup>1</sup>U.S. Naval Research Laboratory

<sup>2</sup>Yale University

<sup>3</sup>University of Minnesota

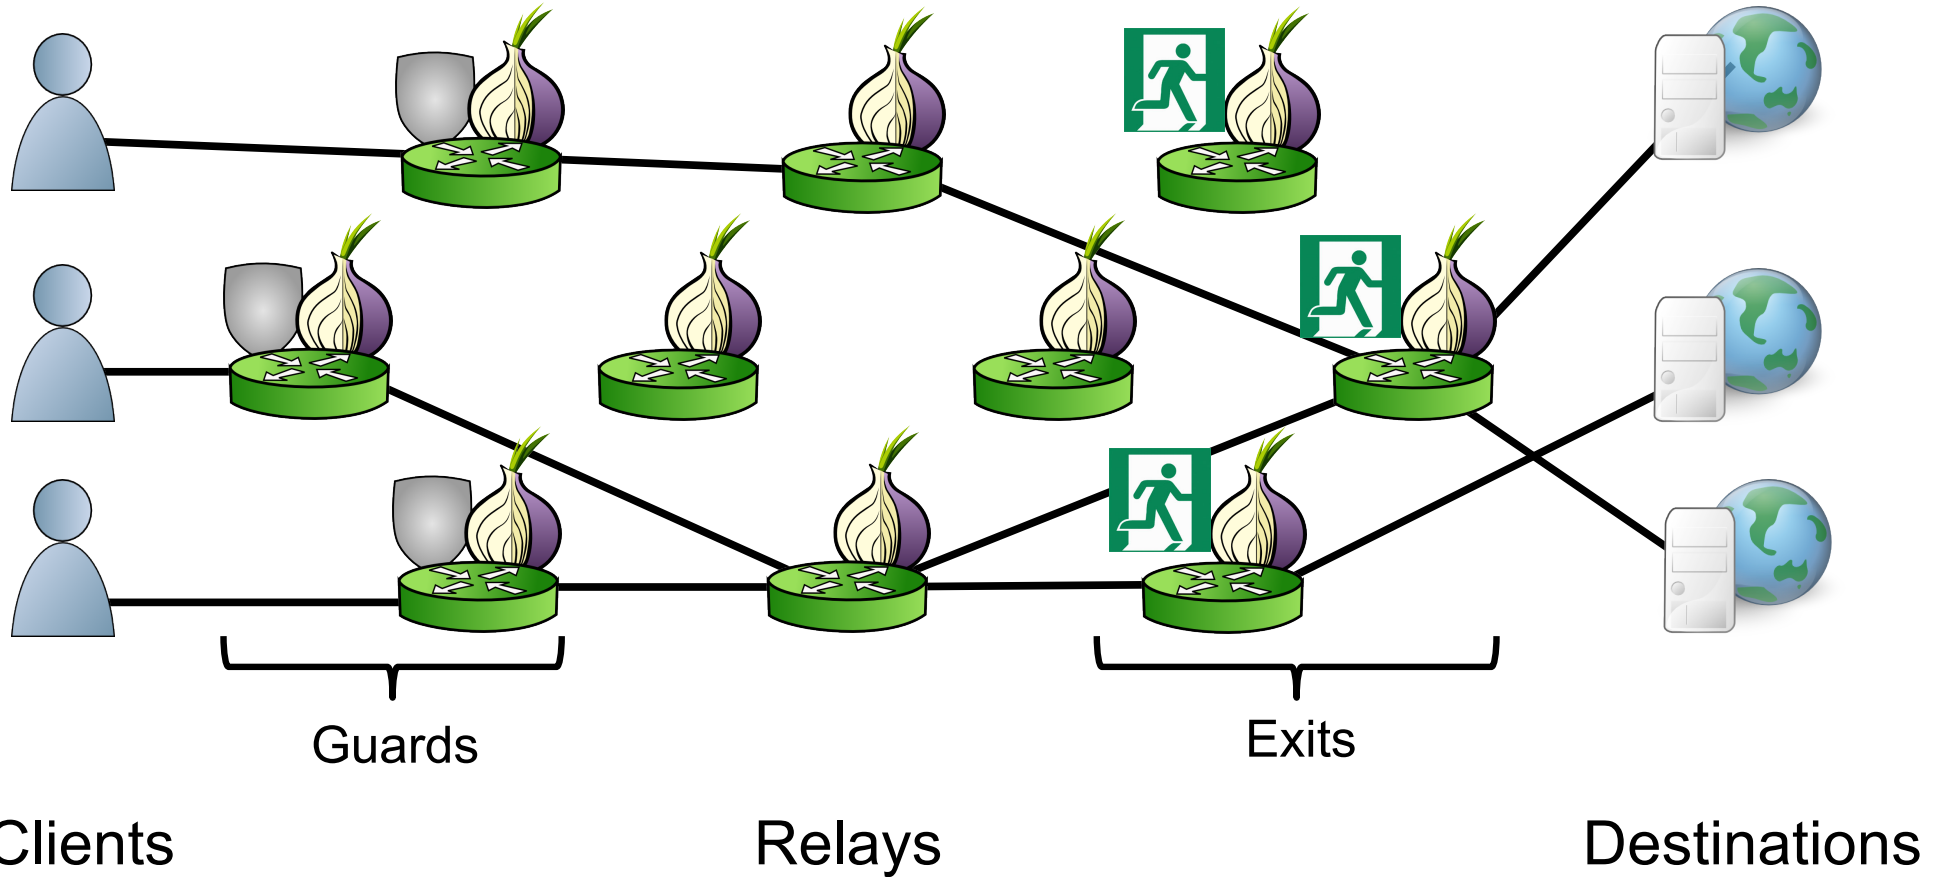
July 18<sup>th</sup>, 2017

Privacy Enhancing Technologies Symposiu

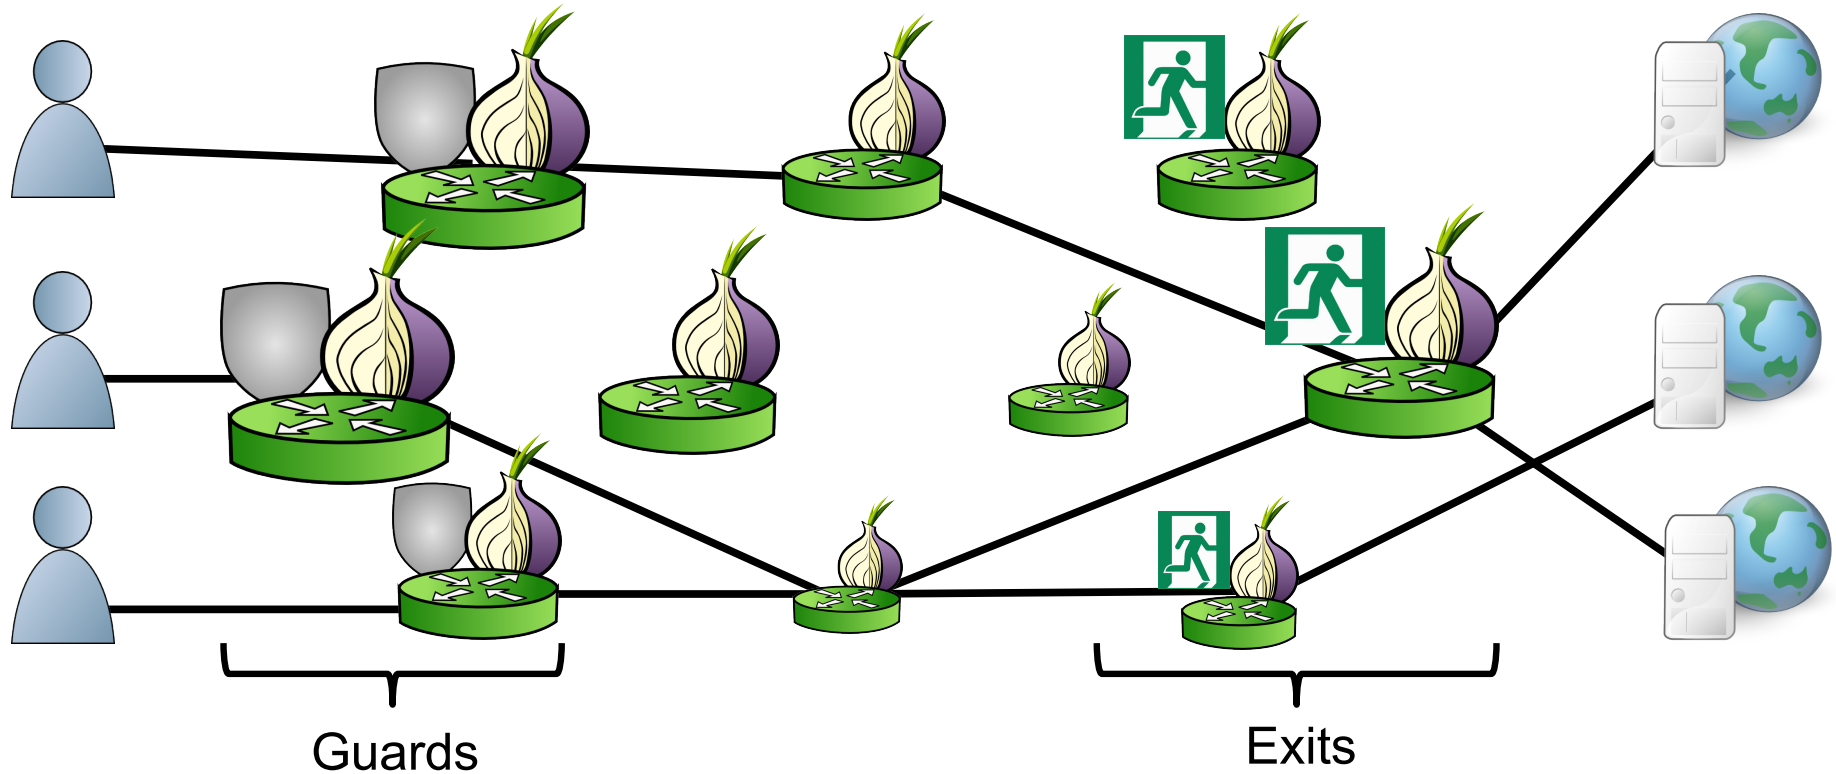
- Problem: Secure load-balancing in Tor
  - Existing Solutions
    - TorFlow
    - EigenSpeed
  - New Solution: PeerFlow
- } Demonstrate attacks
- Prove security against bandwidth-limited adversary
  - Experiments show similar performance to TorFlow

- **Problem: Secure load-balancing in Tor**
  - Existing Solutions
    - TorFlow
    - EigenSpeed
  - New Solution: PeerFlow
- } Demonstrate attacks
- Prove security against bandwidth-limited adversary
  - Experiments show similar performance to TorFlow

# Problem



# Problem




Clients

Relays

Destinations

- Tor relays have varying unknown *capacities*

# Problem

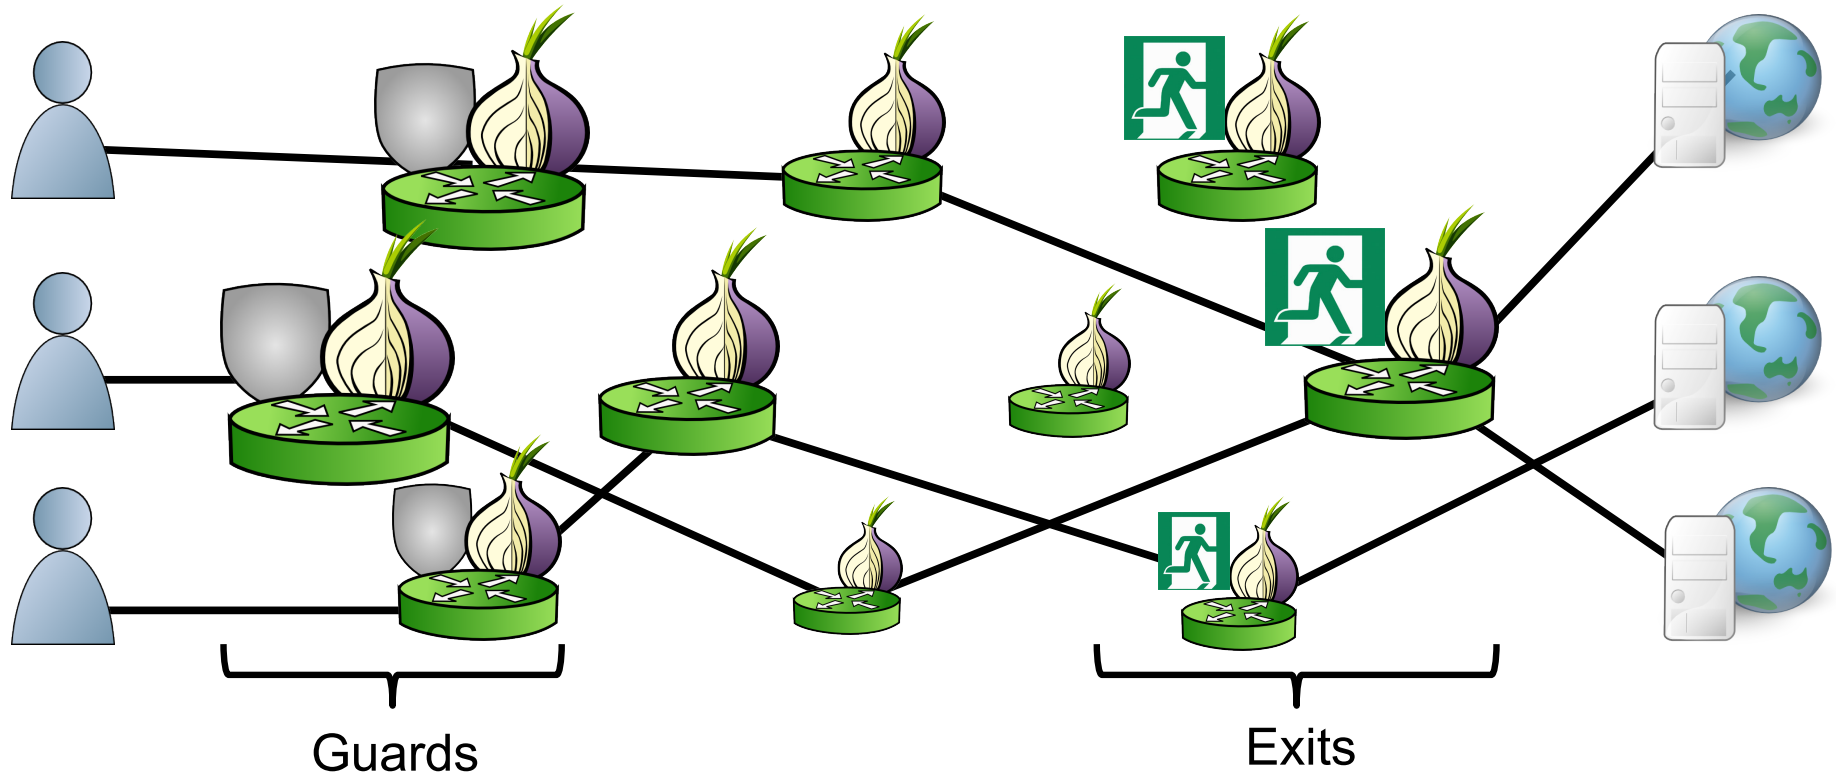


Clients

Relays

Destinations

- Tor relays have varying unknown *capacities*
- Clients must balance load




Clients

Relays

Destinations

- Tor relays have varying unknown *capacities*
- Clients must balance load
- Insecure load balancing allows adversary to attack more client traffic




Clients

Relays

Destinations

- Tor relays have varying unknown *capacities*
- Clients must balance load
- Insecure load balancing allows adversary to attack more client traffic






Clients

Relays

Destinations

- Tor relays have varying unknown *capacities*
- Clients must balance load
- Insecure load balancing allows adversary to attack more client traffic

# Problem



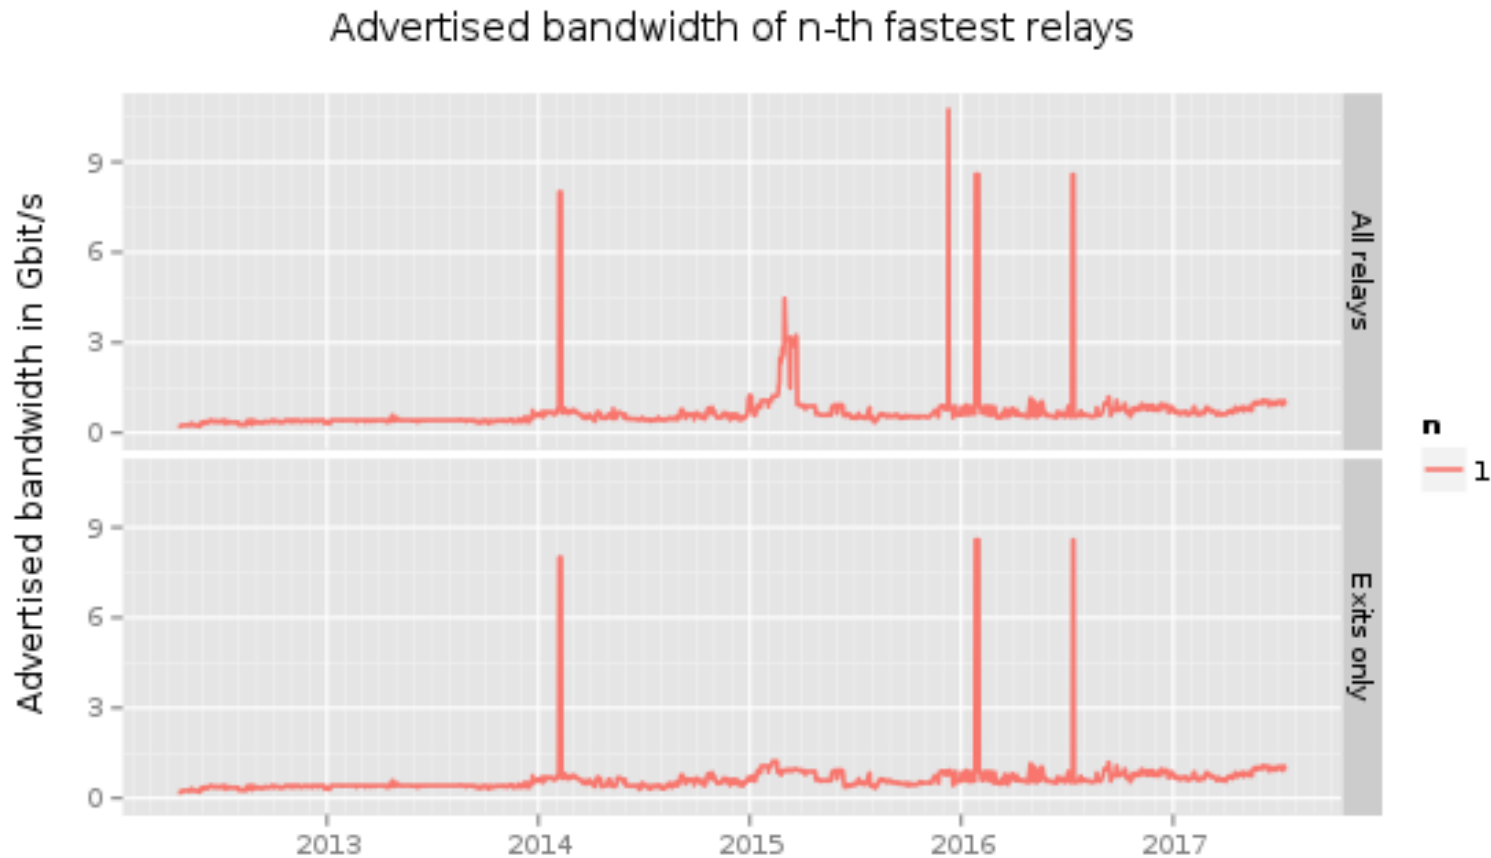
Clients

Relays

Destinations

- Tor relays have varying unknown *capacities*
- Clients must balance load
- Insecure load balancing allows adversary to attack more client traffic

The threat is real: relay falsely advertise bandwidth.




The Tor Project - <https://metrics.torproject.org/>

- **Problem: Secure load-balancing in Tor**
  - Existing Solutions
    - TorFlow
    - EigenSpeed
  - New Solution: PeerFlow
- } Demonstrate attacks
- Prove security against bandwidth-limited adversary
  - Experiments show similar performance to TorFlow

- Problem: Secure load-balancing in Tor
- Existing Solutions
  - TorFlow
  - EigenSpeed } Demonstrate attacks
- New Solution: PeerFlow
  - Prove security against bandwidth-limited adversary
  - Experiments show similar performance to TorFlow

## Design

1. Relays are divided into 50-relay slices by estimated capacity.
2. Bandwidth Authorities (BWAuths) time fetching test files through pairs of relay in each slice.
3. Relays given capacities by multiplying self-reported bandwidth by test speed divided by average speed.

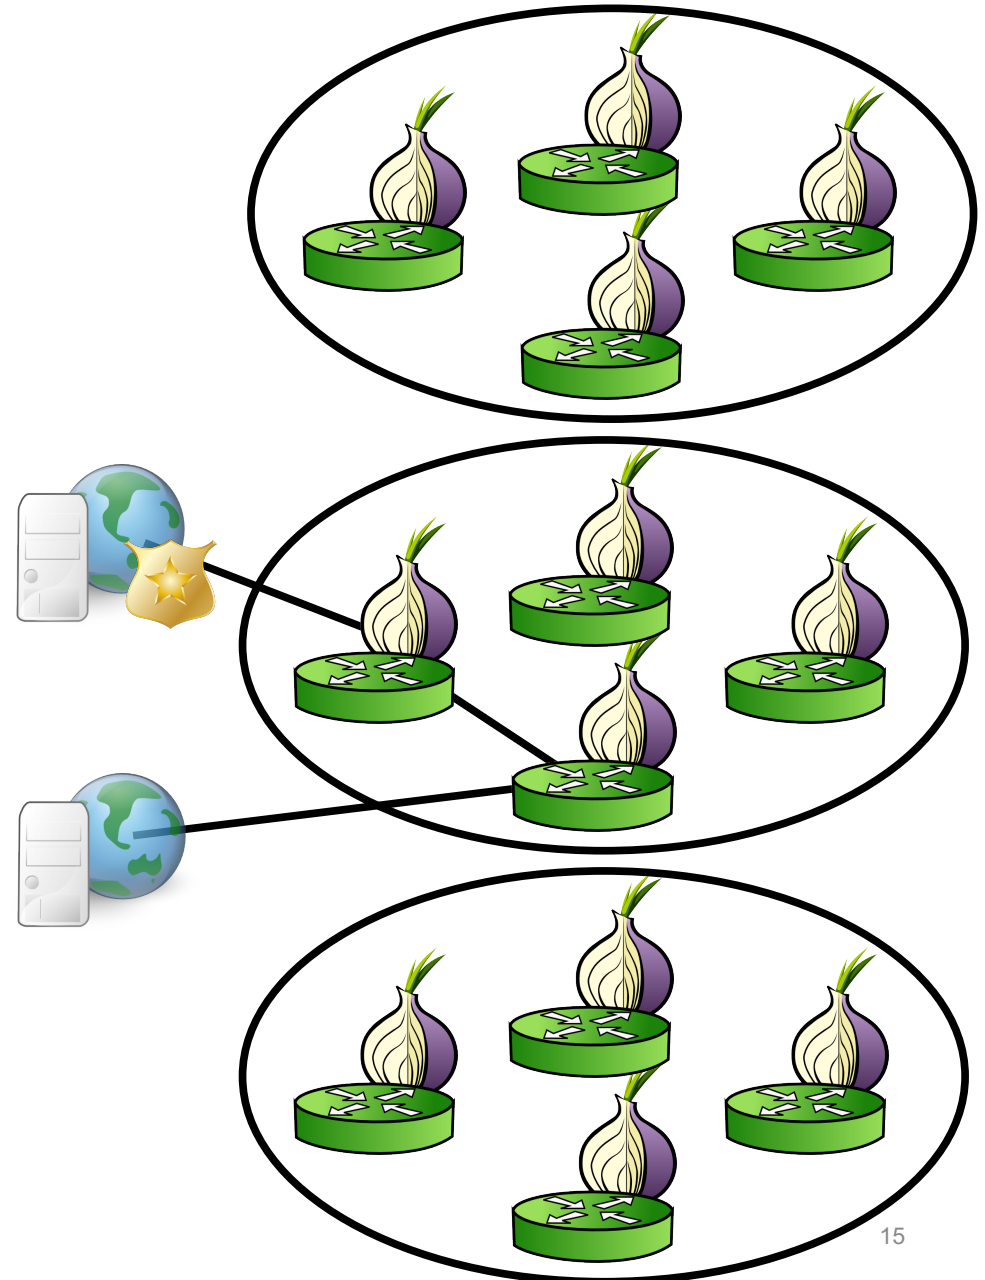


## Design

1. Relays are divided into 50-relay slices by estimated capacity.
2. Bandwidth Authorities (BWAuths) time fetching test files through pairs of relay in each slice.
3. Relays given capacities by multiplying self-reported bandwidth by test speed divided by average speed.

## Attacks

1. Self-reported bandwidth can be set arbitrarily high.
2. Relays can recognize test downloads and relay data only in those cases
3. Malicious pairs need not actually download the file (no validation).

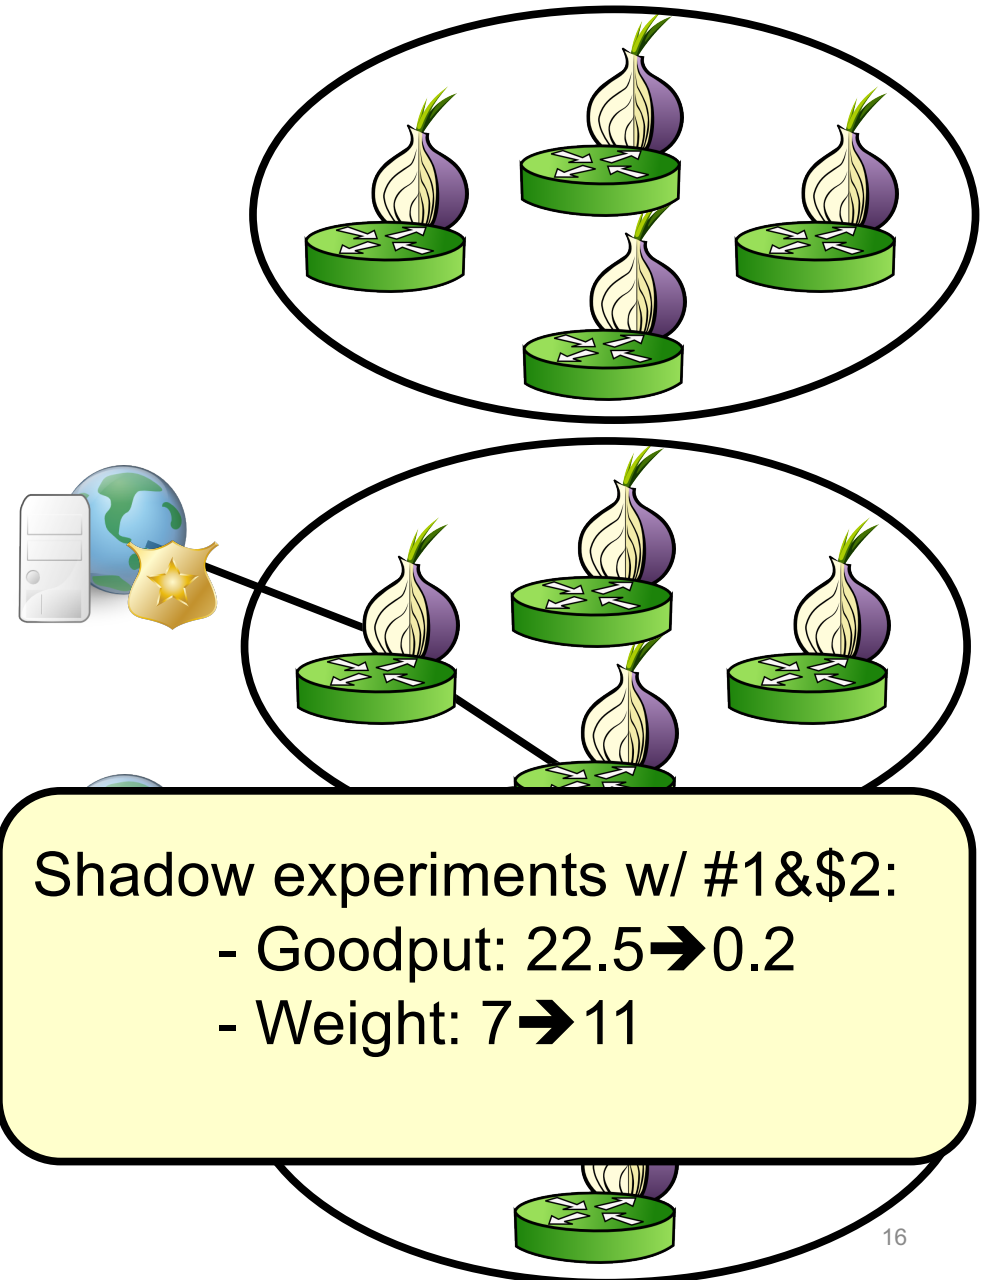


## Design

1. Relays are divided into 50-relay slices by estimated capacity.
2. Bandwidth Authorities (BWAuths) time fetching test files through pairs of relay in each slice.
3. Relays given capacities by multiplying self-reported bandwidth by test speed divided by average speed.

## Attacks

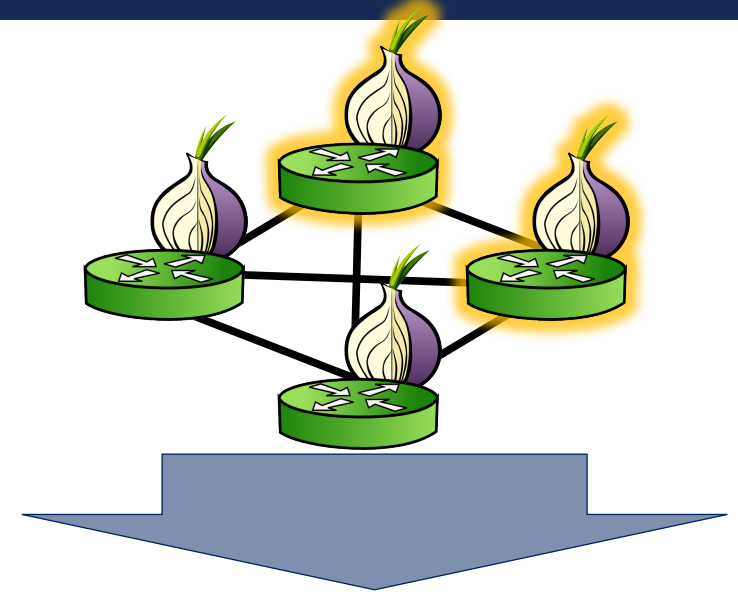
1. Self-reported bandwidth can be set arbitrarily high.
2. Relays can recognize test downloads and relay data only in those cases
3. Malicious pairs need not actually download the file (no validation).





## Design

1. Relays periodically send max speed of other relays to a BWAAuth.
2. Aggregator calculates capacities as eigenvector of largest connected component with *trusted* relays.
3. Exclude as “liars” relays w/ reports
  1. Changing too quickly during computation, or
  2. Too different from eigenvector



$T =$

0	$s_{12}$	$s_{13}$	$s_{14}$
$s_{21}$	0	$s_{23}$	$s_{24}$
$s_{31}$	$s_{32}$	0	$s_{34}$
$s_{41}$	$s_{42}$	$s_{43}$	0




Normalize  $T$ :  $T'$

Output  $v^*$ :  $v^*T = \lambda T$ ,  $\lambda \geq 1$

## Design

1. Relays periodically send max speed of other relays to a BWAuth.
2. Aggregator calculates capacities as eigenvector of largest connected component with *trusted* relays.
3. Exclude as “liars” relays w/ reports
  1. Changing too quickly during computation, or
  2. Too different from eigenvector



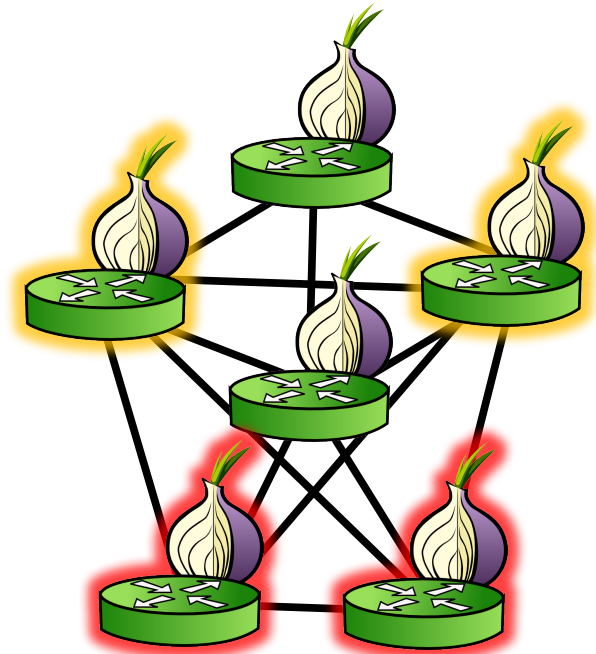
**Fat-pipe attack:** Large false speeds among malicious relays, small elsewhere. EigenSpeed’s liar detection is designed to prevent this.

## Design

1. Relays periodically send max speed of other relays to a BWAuth.
2. Aggregator calculates capacities as eigenvector of largest connected component with *trusted* relays.
3. Exclude as “liars” relays w/ reports
  1. Changing too quickly during computation, or
  2. Too different from eigenvector

## Attack

1. “Frame” some honest non-trusted relays under liar metric #1 with avg speeds with all but framed relays.



**Framing attack:** With 1118 trusted relays and 2.83% malicious BW, and 558 malicious relays, 559 of 5000 honest relays are framed.


- Problem: Secure load-balancing in Tor
- Existing Solutions
  - TorFlow
  - EigenSpeed
- **New Solution: PeerFlow**
  - Prove security against bandwidth-limited adversary
  - Experiments show similar performance to TorFlow



# PeerFlow: Design


# PeerFlow: Design

1. *Measuring relays* (largest by capacity) record total bytes transferred with all other relays.



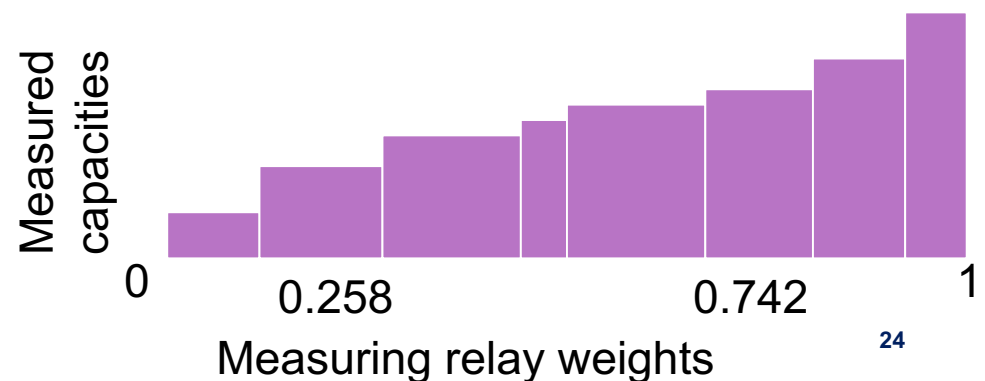

# PeerFlow: Design

1. *Measuring relays* (largest by capacity) record total bytes transferred with all other relays
2. Measurements added to random noise and divided by position probabilities. Result ( $\rho_i$ ) submitted to BW Authorities (BWAAuths).



# PeerFlow: Design

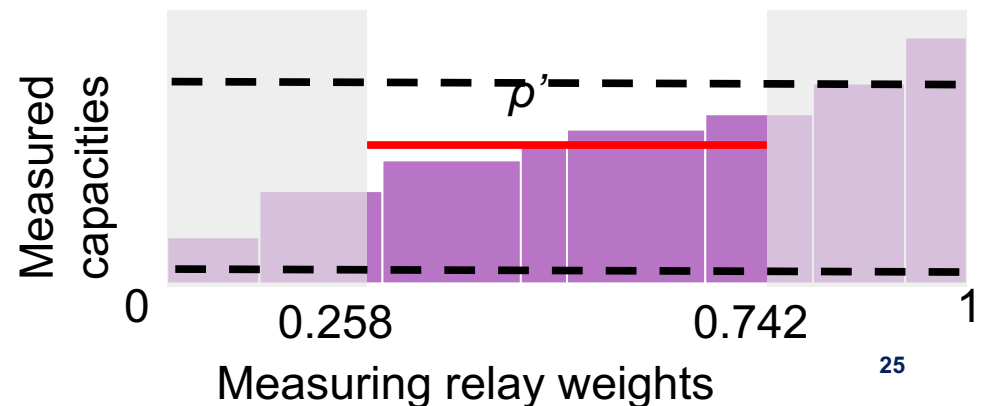

1. *Measuring relays* (largest by capacity) record total bytes transferred with all other relays
2. Measurements added to random noise and divided by position probabilities. Result ( $\rho_i$ ) submitted to BW Authorities (BWAAuths).
3. BWAAuths estimate the total bytes relayed  $\rho'$  as the windowed, trimmed mean, trimming fractions by current capacity and windowing from trusted measurements.






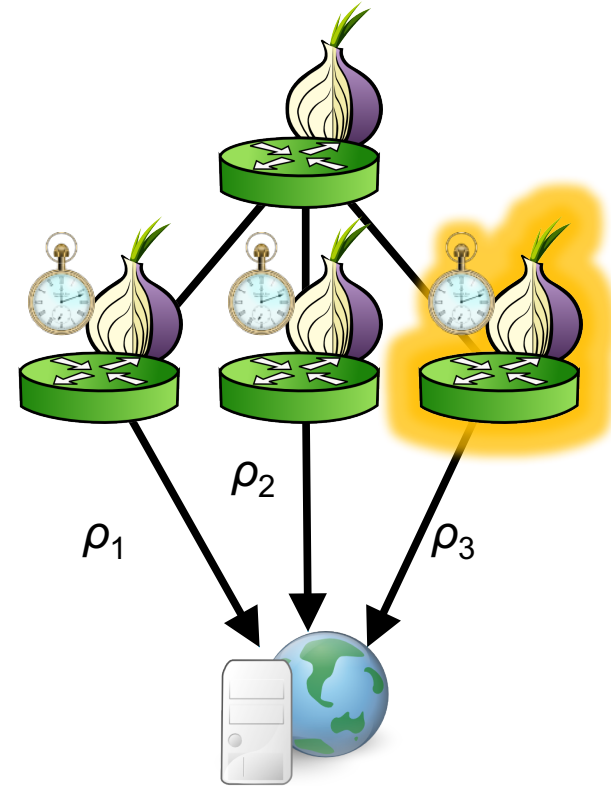
# PeerFlow: Design

1. *Measuring relays* (largest by capacity) record total bytes transferred with all other relays
2. Measurements added to random noise and divided by position probabilities. Result ( $\rho_i$ ) submitted to BW Authorities (BWAAuths).
3. BWAAuths estimate the total bytes relayed  $\rho'$  as the windowed, trimmed mean, trimming fractions by current capacity and windowing from trusted measurements.





# PeerFlow: Design

1. *Measuring relays* (largest by capacity) record total bytes transferred with all other relays
2. Measurements added to random noise and divided by position probabilities. Result ( $\rho_i$ ) submitted to BW Authorities (BWAAuths).
3. BWAAuths estimate the total bytes relayed  $\rho'$  as the windowed, trimmed mean, trimming fractions by current capacity and windowing from trusted measurements.
4. If  $\rho'$  is comparable to that of peers, capacity updated using  $\rho'$ , else relay enters probation.



# PeerFlow: Design

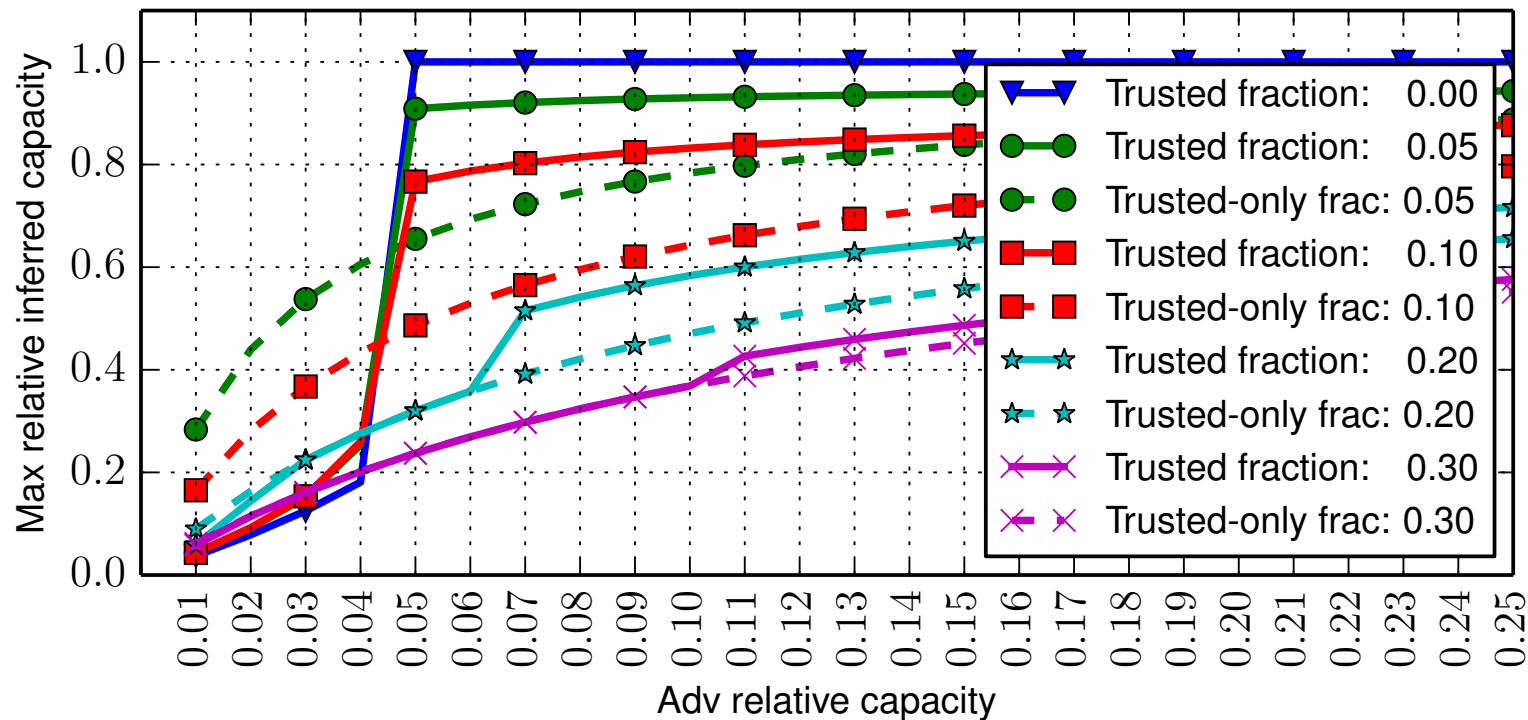
1. *Measuring relays* (largest by capacity) record total bytes transferred with all other relays
2. Measurements added to random noise and divided by position probabilities. Result ( $\rho_i$ ) submitted to BW Authorities (BWAAuths).
3. BWAAuths estimate the total bytes relayed  $\rho'$  as the windowed, trimmed mean, trimming fractions by current capacity and windowing from trusted measurements.
4. If  $\rho'$  is comparable to that of peers, capacity updated using  $\rho'$ , else relay enters probation.
5. New relays only selected for **middle position**



# PeerFlow: Security

Attack	Weight multiple
Only carry traffic in one direction	2
Only exchange traffic with measuring relays	1.33
Do not exchange traffic with the lower trimmed fraction of relays	1.34

## Single-round capacity inflation

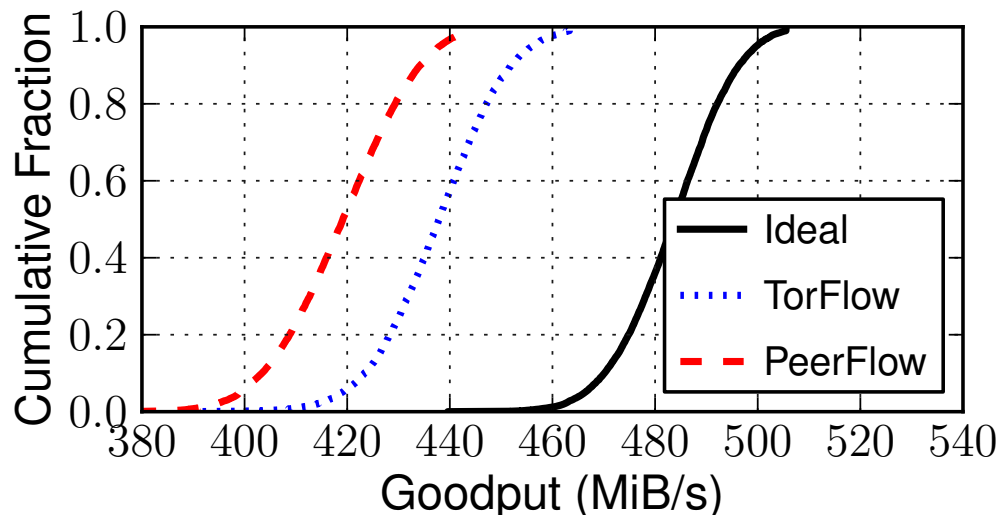


## Multiple-round capacity inflation


# PeerFlow: Performance

Shadow experiments comparing PeerFlow, TorFlow, and Ideal

- 4 Tor directory authorities
- 498 Tor relays
- 7,500 Tor clients
- 1,000 servers



Aggregate relay goodput per second



Time to last byte of 320KiB file

1. Tor needs secure load balancing
2. Demonstrated attacks on existing solutions
  - TorFlow
  - EigenSpeed
3. Presented PeerFlow
  - Demonstrated secure against bandwidth-limited adversary
  - Experimentally showed performance is similar to current Tor performance