ExperimenTor: A Testbed for Safe and Realistic Tor Experimentation

Kevin Bauer[1] Micah Sherr[2]
Damon McCoy[3] Dirk Grunwald[4]

[1]University of Waterloo [2]Georgetown University
[3]UCSD [4]University of Colorado

`http://crysp.uwaterloo.ca/software/exptor`

4th Workshop on Cyber Security Experimentation and Test (CSET '11)

What is Tor and why is it important?

Tor is a *low-latency* overlay network and a software package that allows you to use TCP-based applications *anonymously*

Tor has an estimated 350,000 daily users world-wide and its network consists of over 2,500 volunteer-operated Tor routers

Ordinary Citizens

- Protect web browsing habits
- Research sensitive or taboo topics
- Circumvent censorship

Activists & Whistleblowers

- Expose human rights violations
- Promote democracy
- Protest election results

Corporations

- Research the competition
- Safeguard trade secrets

Law Enforcement

- Online surveillance
- Sting operations

Tor uses layered encryption to hide your online behaviors

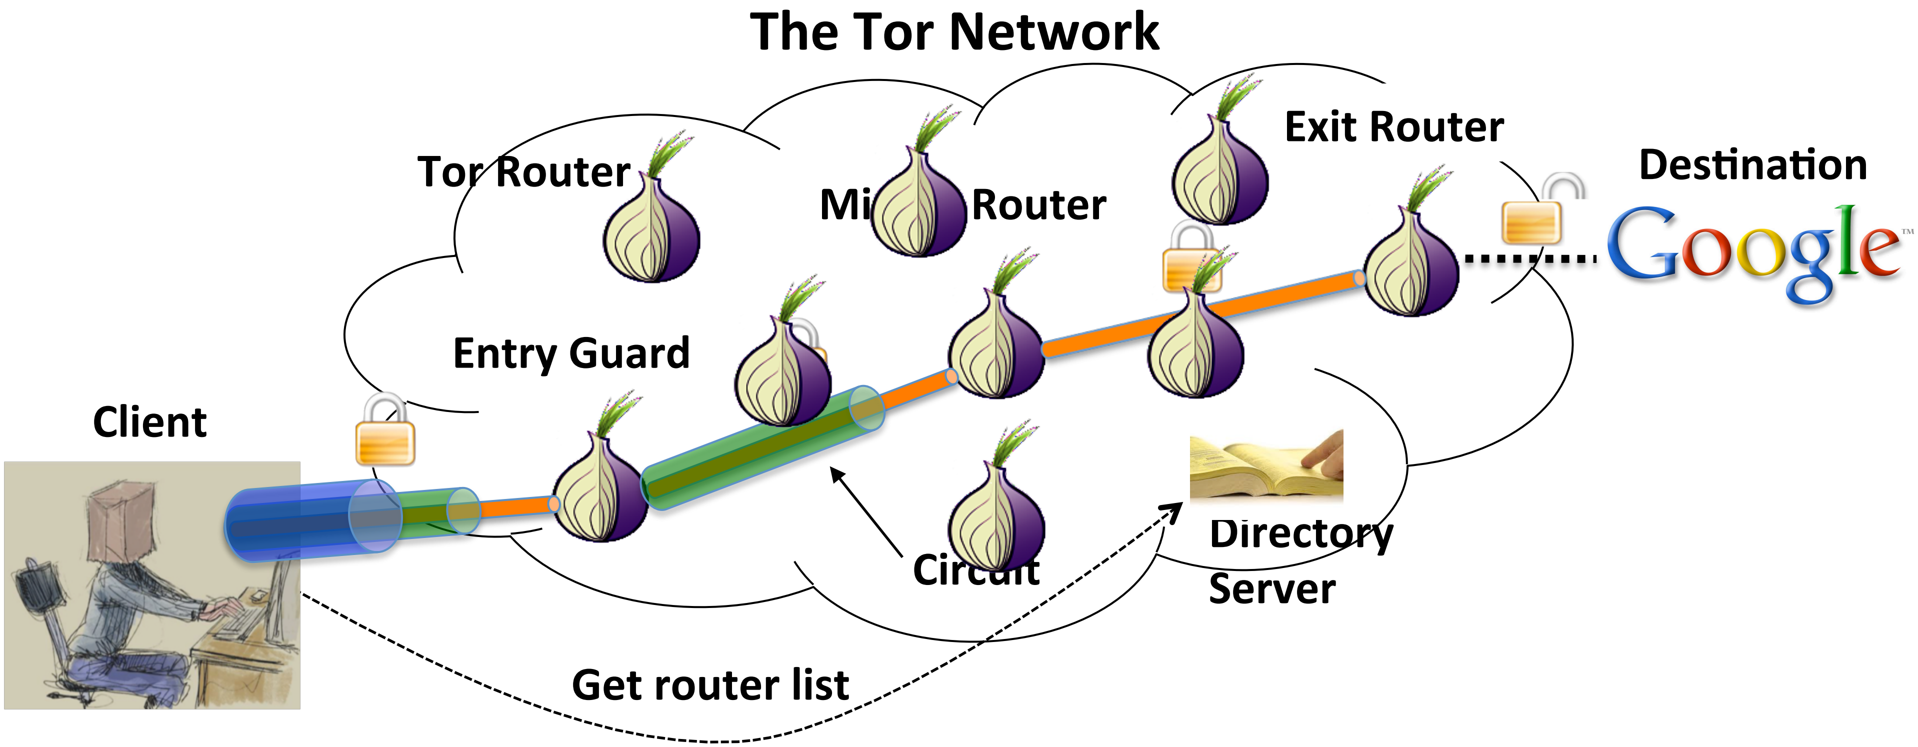

Tor provides *anonymity for TCP applications* by tunneling traffic through a *virtual circuit* of three Tor routers using layered encryption

Communicating parties are *unlinkable* as long as the entry and exit routers do not collude

Tor is still an evolving research network

- Past and current research aims to improve Tor's:
  - Security and anonymity [CCS '07, NDSS '08, USENIX Security '10]
  - Quality of service [USENIX Security '09, CCS '10, PETS '11]
- **Problem:** There is no standard methodology for conducting Tor research in a *realistic* and *safe* manner; prior methods include:

Realism →

Abstract modeling

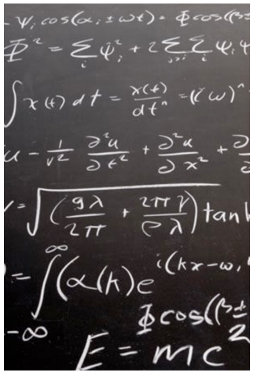

Network simulators

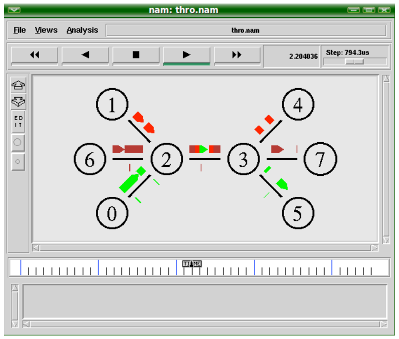

Small-scale network emulation

emulab

deterlab
based on emulab

Distributed research networks

PLANETLAB

Live Tor network

Tor

← **Safety**

ExperimenTor: A whole-network Tor emulation testbed

Goal: Propose a standard experimental methodology

ExperimenTor

- Replicates all components of the Tor network *in isolation*
- Reproduces plausible network conditions through *scalable network emulation*
- Fuels experiments with *empirically derived models*

Allows investigators to study global, whole-network effects

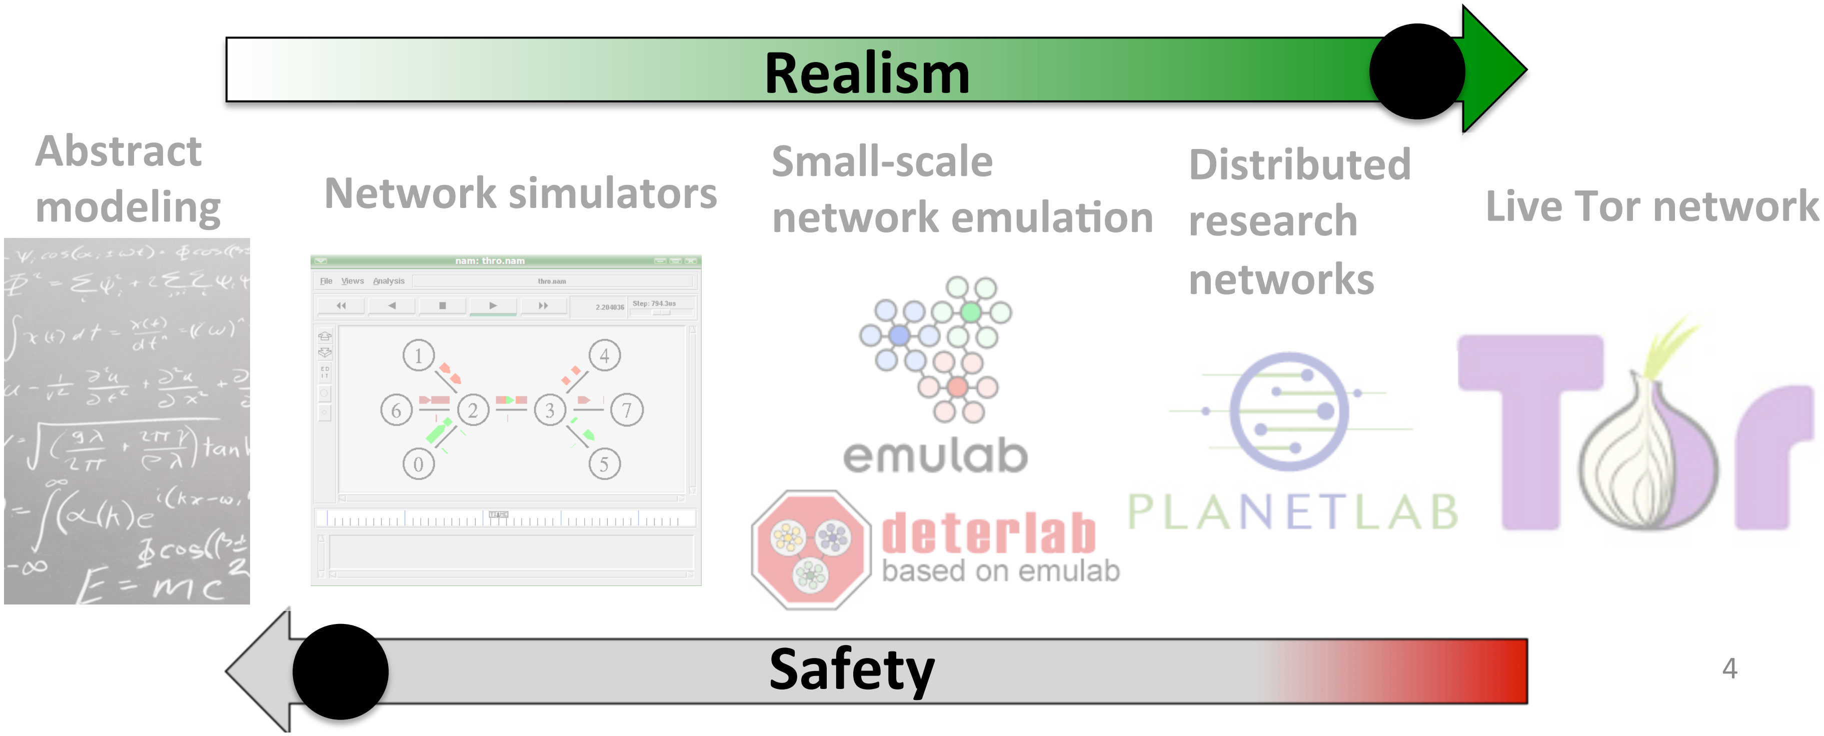

Realism

Abstract modeling

Network simulators

Small-scale network emulation

Distributed research networks

Live Tor network

Safety

Talk outline

- Motivating case studies from prior Tor research
- Challenges of building a Tor network testbed
- Design and implementation of *ExperimenTor*
- Early experiences and lessons learned
- Conclusions and future work

Case study: Whole-network PlanetLab experiments

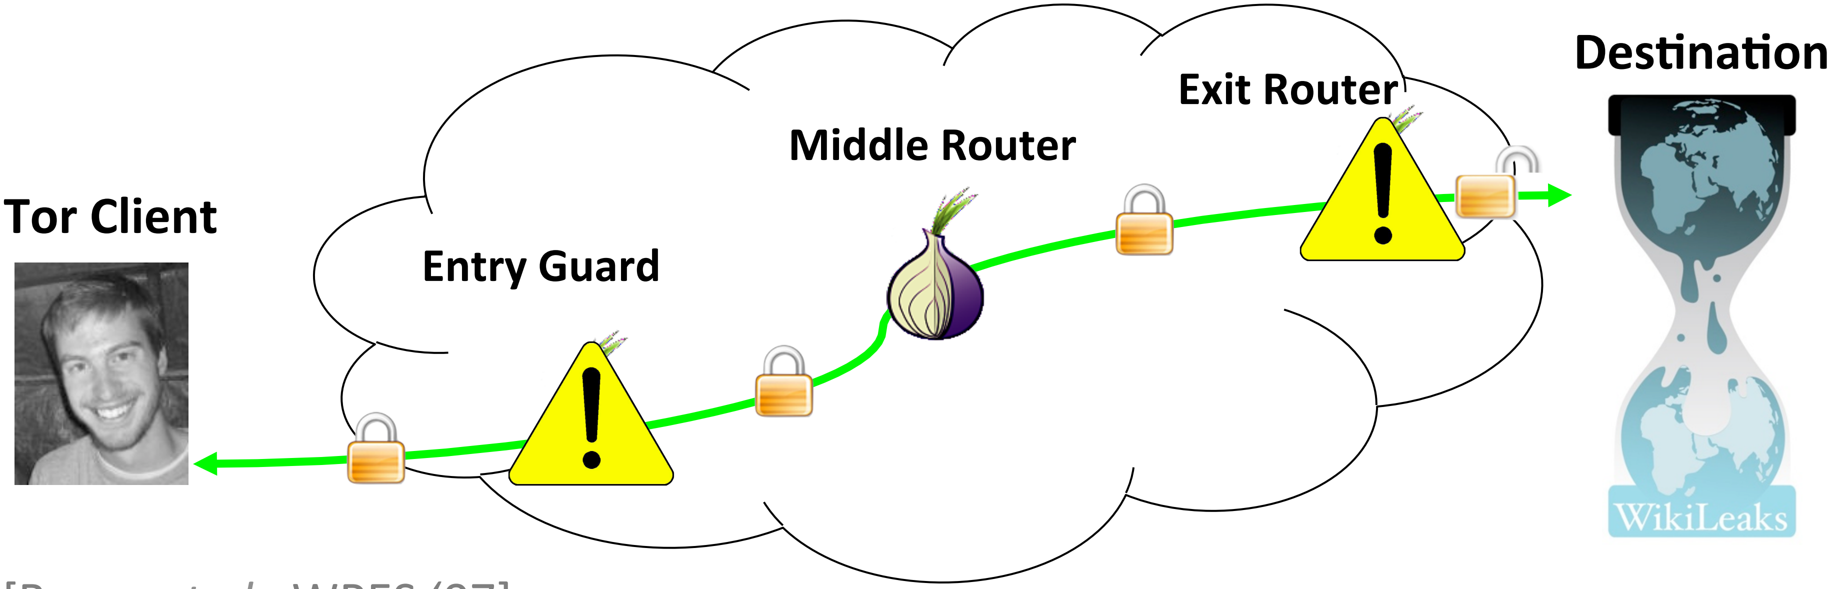

[Bauer *et al.*, WPES '07]

It is assumed that an attacker who controls an entry/exit pair can trivially link the communicating parties: **traffic confirmation attack**

Uniform router selection: Probability of attack's success is $(c/n)^2$, c malicious routers in a network of n total routers

Tor routers are selected in proportion to their perceived bandwidth capacities for load balancing, *but malicious routers can lie*

Case study: Whole-network PlanetLab experiments (2)

Experiment: Evaluated the attack on two small Planetlab deployments with 40 and 60 honest Tor routers

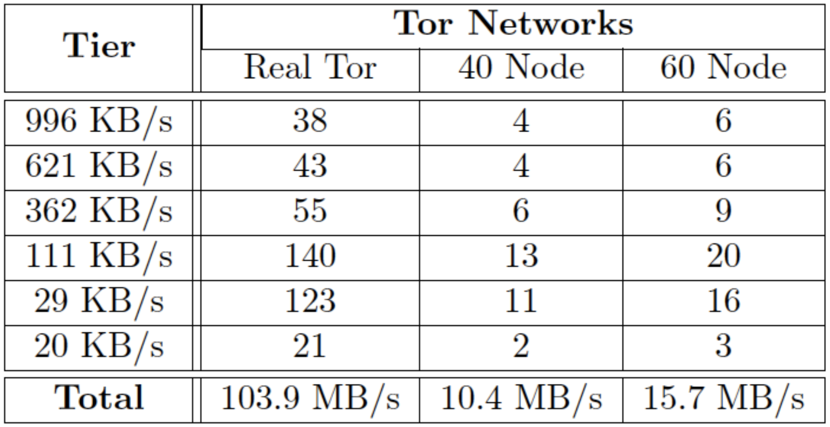

Tier	Tor Networks		
	Real Tor	40 Node	60 Node
996 KB/s	38	4	6
621 KB/s	43	4	6
362 KB/s	55	6	9
111 KB/s	140	13	20
29 KB/s	123	11	16
20 KB/s	21	2	3
Total	103.9 MB/s	10.4 MB/s	15.7 MB/s

Details: Sample the bandwidth distribution of the real Tor network

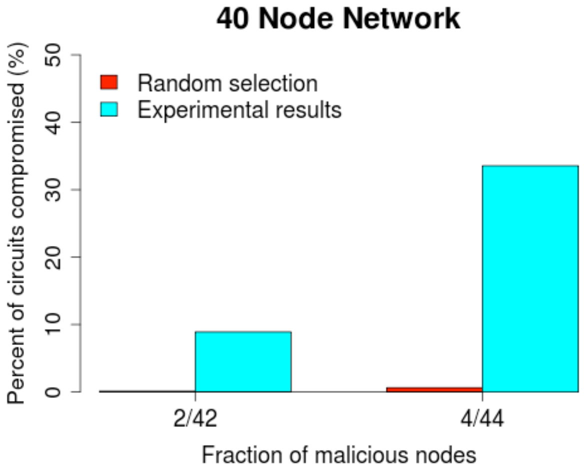

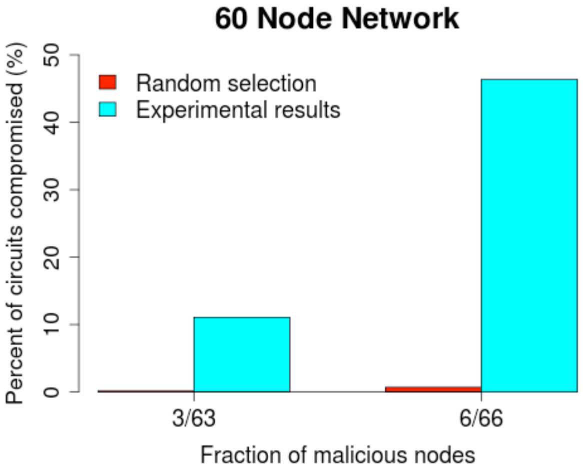

[Bauer *et al.*, WPES '07]

Limitations:

1. Reduced scale
2. Need to run many measurements to find suitable PlanetLab nodes
3. Repeatability?

Case study: Small-scale experiments with the live Tor network

Tunable Tor was proposed to help users manage their risk of the previous attack [Snader and Borisov, NDSS ’08]

Uniform selection **Skewed to high bandwidth nodes**

User-tunable router selection

High anonymity **High performance**

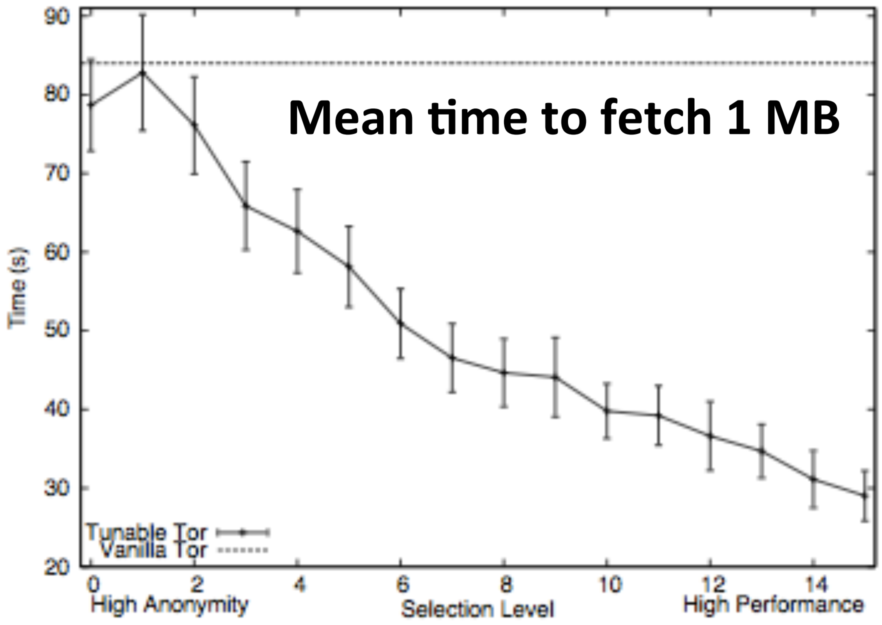

Mean time to fetch 1 MB

Experiment : Deployed one “Tunable Tor” client on the live Tor network

Details: Measured download times at different “selection levels”

Limitations: What happens when *many* Tor clients use Tunable Tor? Global effects?

A case for whole-network Tor emulation

Goal: Capture all salient dynamics of the live Tor network and reproduce in isolation → Realistic and safe experiments

Desired features:

1. Allow investigators to deploy small-scale, large-scale, or global changes to any part of Tor's design
2. Should eliminate any risk to the live Tor network
3. Experimental results should be meaningful to the live Tor network

Our argument: All can be realized with whole-network Tor emulation

Design challenges

- Modeling the live Tor network is difficult
  - **Tor routers:** Bandwidths, guard statuses, exit policies?
  - **Tor clients:** How many? Applications? Behaviors?
- Large-scale network emulation
  - **Emulab** and **DETER** have **limited** and **shared** resources
- Need to run unmodified Tor and application code
  - Avoids re-implementation errors; promotes realism

Meeting the design challenges

Modeling Tor routers:

- Publicly-available router metadata from Tor's directories
- Historical router data aggregated by the Tor Metrics Portal

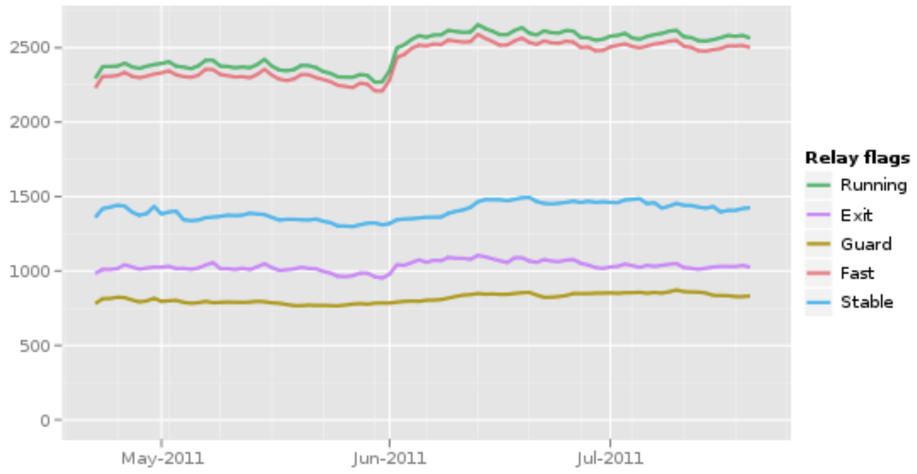

Number of relays with relay flags assigned

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

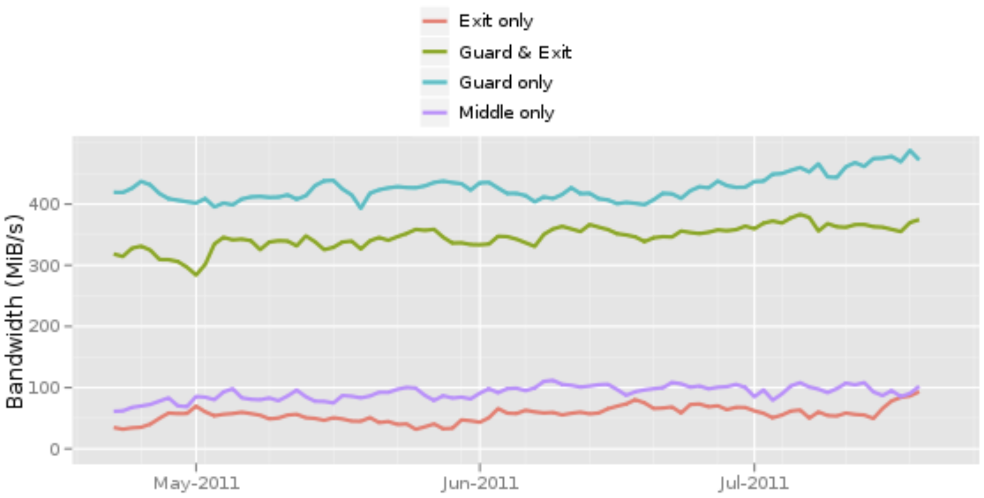

Bandwidth history by relay flags

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

Replicate live Tor's router state, or scale things up or down

Meeting the design challenges (2)

Modeling Tor clients: Leverage existing empirical data on Tor clients and their behaviors [McCoy *et al.*, PETS '08]

Number of Exit TCP Connections by Protocol

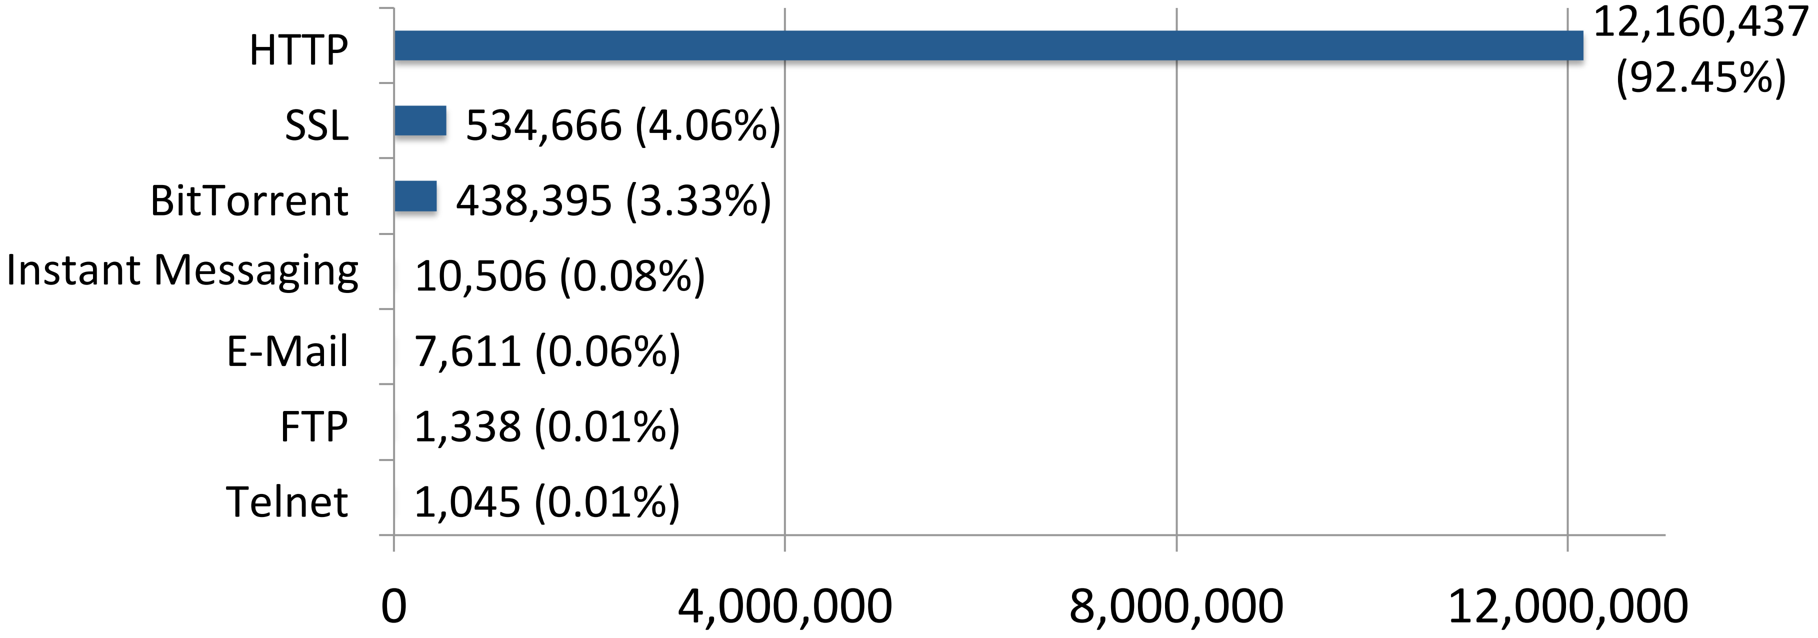

Also leverage existing empirical studies of HTTP traffic to emulate realistic workloads

[*e.g.,* Hernándes-Campos *et al.*, MASCOTS '03; Google web metrics 2010]

Meeting the design challenges (3)

Modeling Tor clients: Leverage existing empirical data on Tor clients and their behaviors [McCoy *et al.*, PETS '08]

Aggregate Exit Traffic Volume by Protocol (GB)

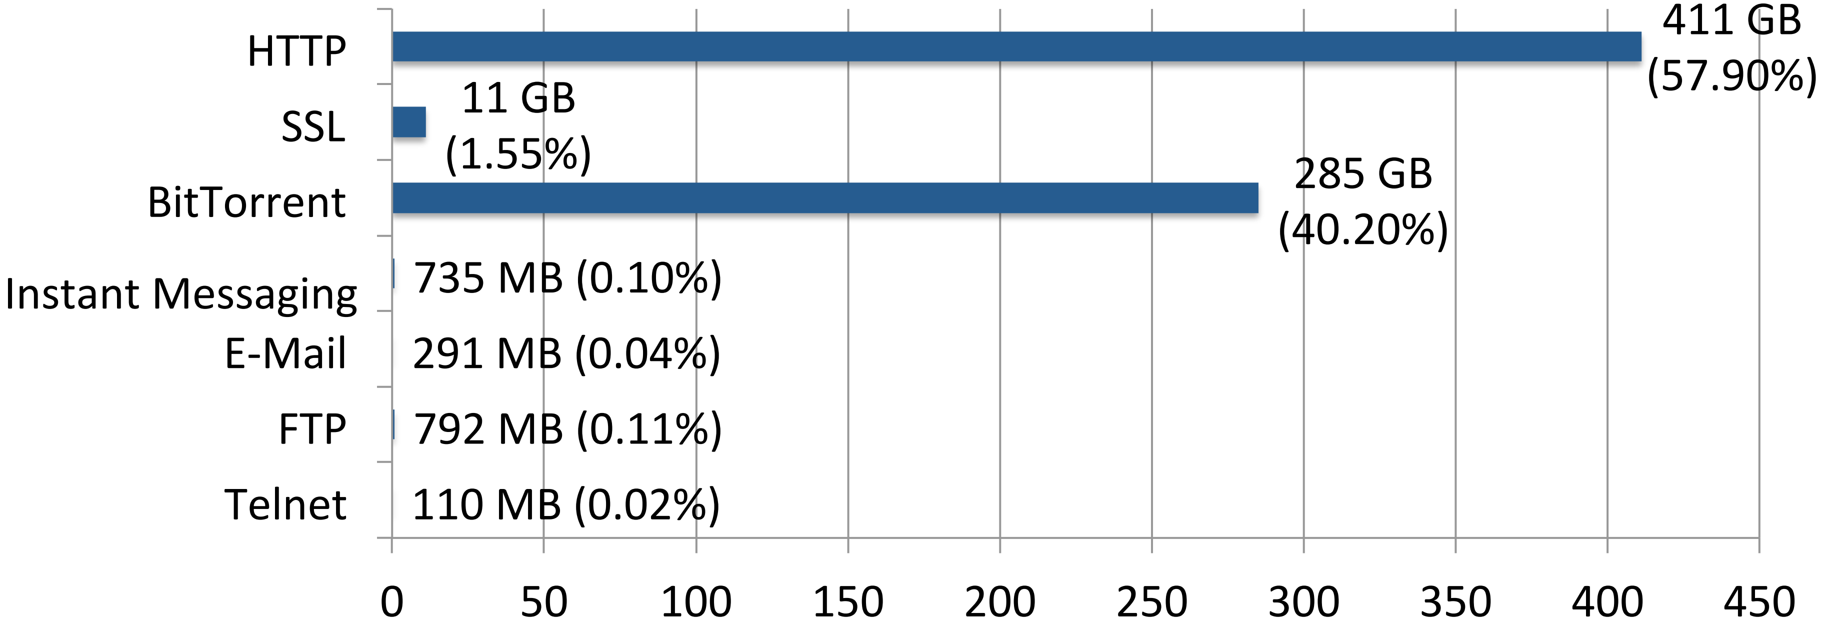

Model the distribution of client traffic by connection and volume

Meeting the design challenges (4)

Modeling Tor clients: Can also leverage publicly-available data about Tor clients from the Tor Metrics Project

Directly connecting users from all countries

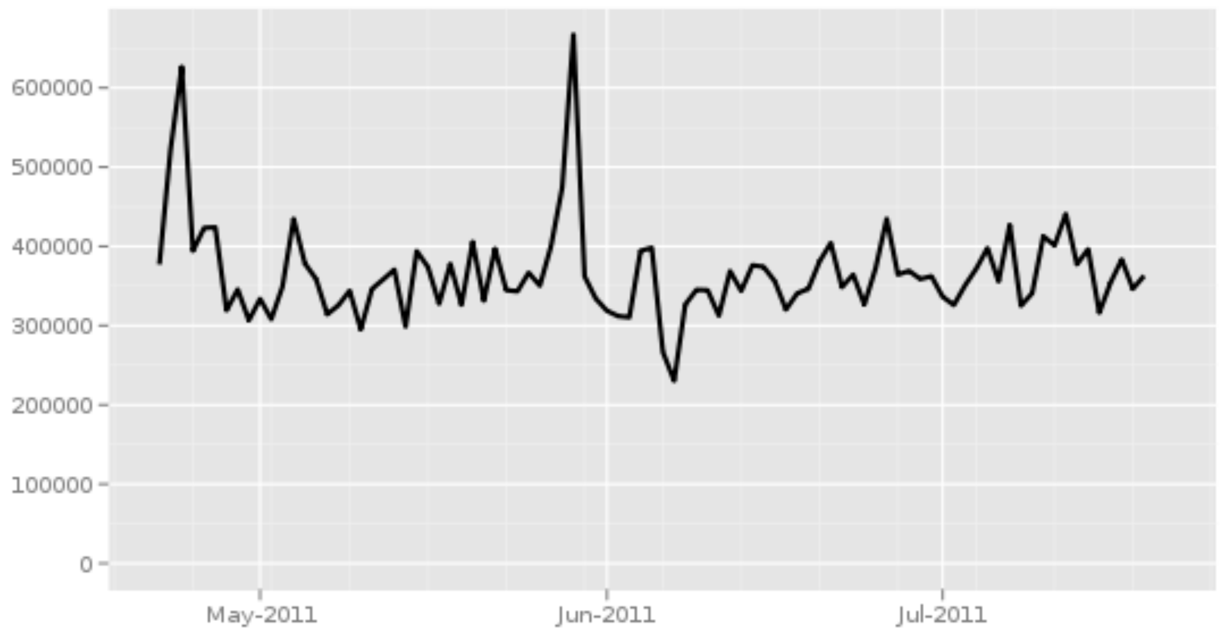

The Tor Project - https://metrics.torproject.org/

Replicate live Tor’s client state, or scale things up or down

Meeting the design challenges (5)

Large-scale network emulation with ModelNet [Vahdat *et al.*, OSDI ‘02]

- Emulates a specified network topology
- Runs native code without modification
- Commodity hardware and OSes; can be deployed at local institution

- High-level system architecture
  - “Emulator” machine: Emulates a network topology in a kernel module
  - “Virtual node” machine: Runs applications within the virtual topology

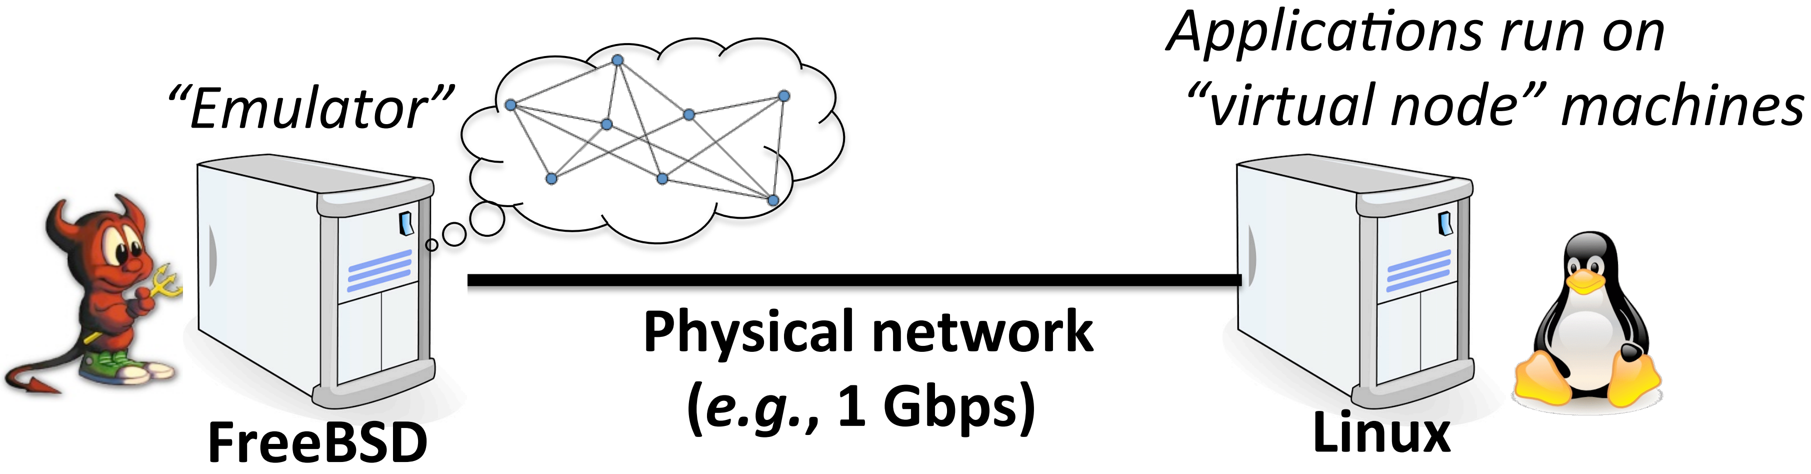

“Emulator”

Applications run on “virtual node” machines

FreeBSD

Physical network (*e.g.*, 1 Gbps)

Linux

Putting it all together

Network topology emulation on a ModelNet core

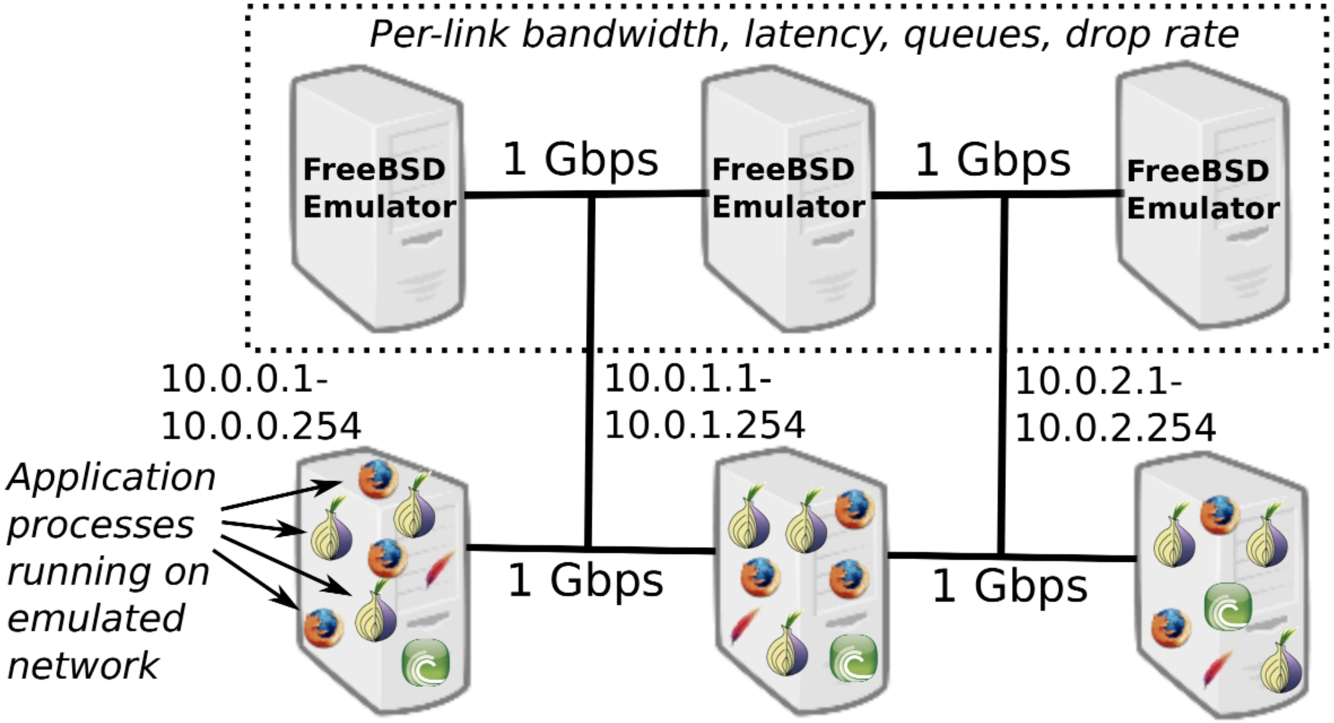

Tor and applications run on edge nodes in virtual topology

Prototype: FreeBSD 6.3
Linux 2.6.32

Accompanying toolkit:

- **Topology generation**
- **Configure Tor clients, routers, & directories**
- **Run experiments & perform analyses**

Testbed and toolkit are publicly available

`http://crysp.uwaterloo.ca/software/exptor`

Early experiences

- ExperimenTor prototypes are deployed at four research institutions (single emulator)
- Used to support two ongoing research projects:
  - Evaluate the effects of link-based router selection
  - Re-design Tor's congestion control and flow control
- Both projects require global design changes to Tor

Limitations and future work

- **Scalability**
  - Scaling experiments to Tor’s estimated 350K users is likely not possible; necessary to “down sample”
- **Improve client and traffic models**
  - Data on Tor usage are limited
  - Is it possible to emulate diverse versions and configurations of Tor users?

Summary and conclusion

- *ExperimenTor* is a whole-network emulation-based testbed and toolkit for safe and realistic Tor experiments
- Enables large-scale Tor experiments that:
  - Use real Tor router bandwidths to inform topology
  - Emulate Tor clients and their traffic
  - Enable experiments with global changes to Tor's design
  - Can be deployed cheaply on commodity systems

For more information:

`http://crysp.uwaterloo.ca/software/exptor`