TECHNISCHE UNIVERSITÄT DARMSTADT

Tor Experimentation Tools

Fatemeh Shirazi

TU Darmstadt / KU Leuven
Darmstadt, Germany
fshirazi@cdc.informatik.tu-darmstadt.de

Matthias Göhring

TU Darmstadt
Darmstadt, Germany
de.m.goehring@ieee.org

Claudia Diaz

KU Leuven / iMinds
Leuven, Belgium
claudia.diaz@esat.kuleuven.be

Tor Experimentation Tools

- Background
- Network Statistics
- How it works
- CollecTor
- Research

Tor Basics

TECHNISCHE UNIVERSITÄT DARMSTADT

- Distributed overlay anonymity network
- Operated by volunteers around the world
- Developed and maintained by The Tor Project (non-profit)
- Active research community

Network Components

TECHNISCHE UNIVERSITÄT DARMSTADT

- Relays: Onion Router (OR)
  - Entry guard
  - Middle node
  - Exit node

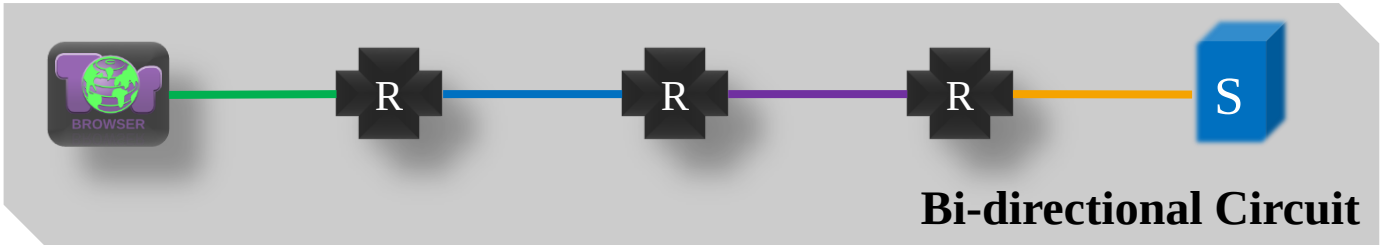

- Client Software: Onion Proxy (OP)
- Directory Servers (Authorities and Mirrors)
- Bridges („*hidden*“ relays)

Tor Network Size

TECHNISCHE UNIVERSITÄT DARMSTADT

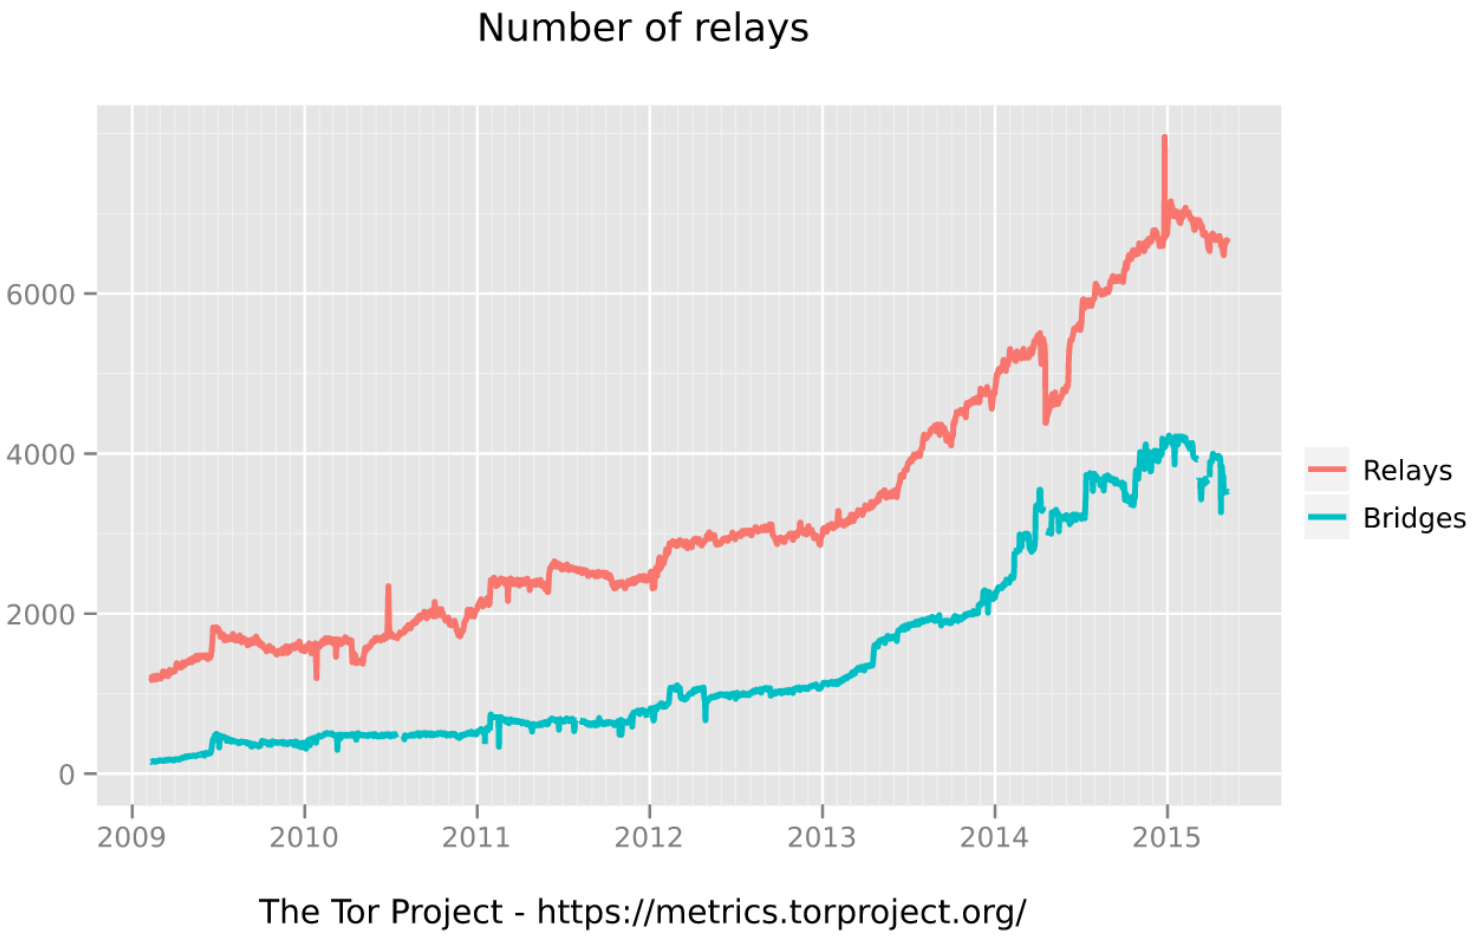

Tor Users

TECHNISCHE UNIVERSITÄT DARMSTADT

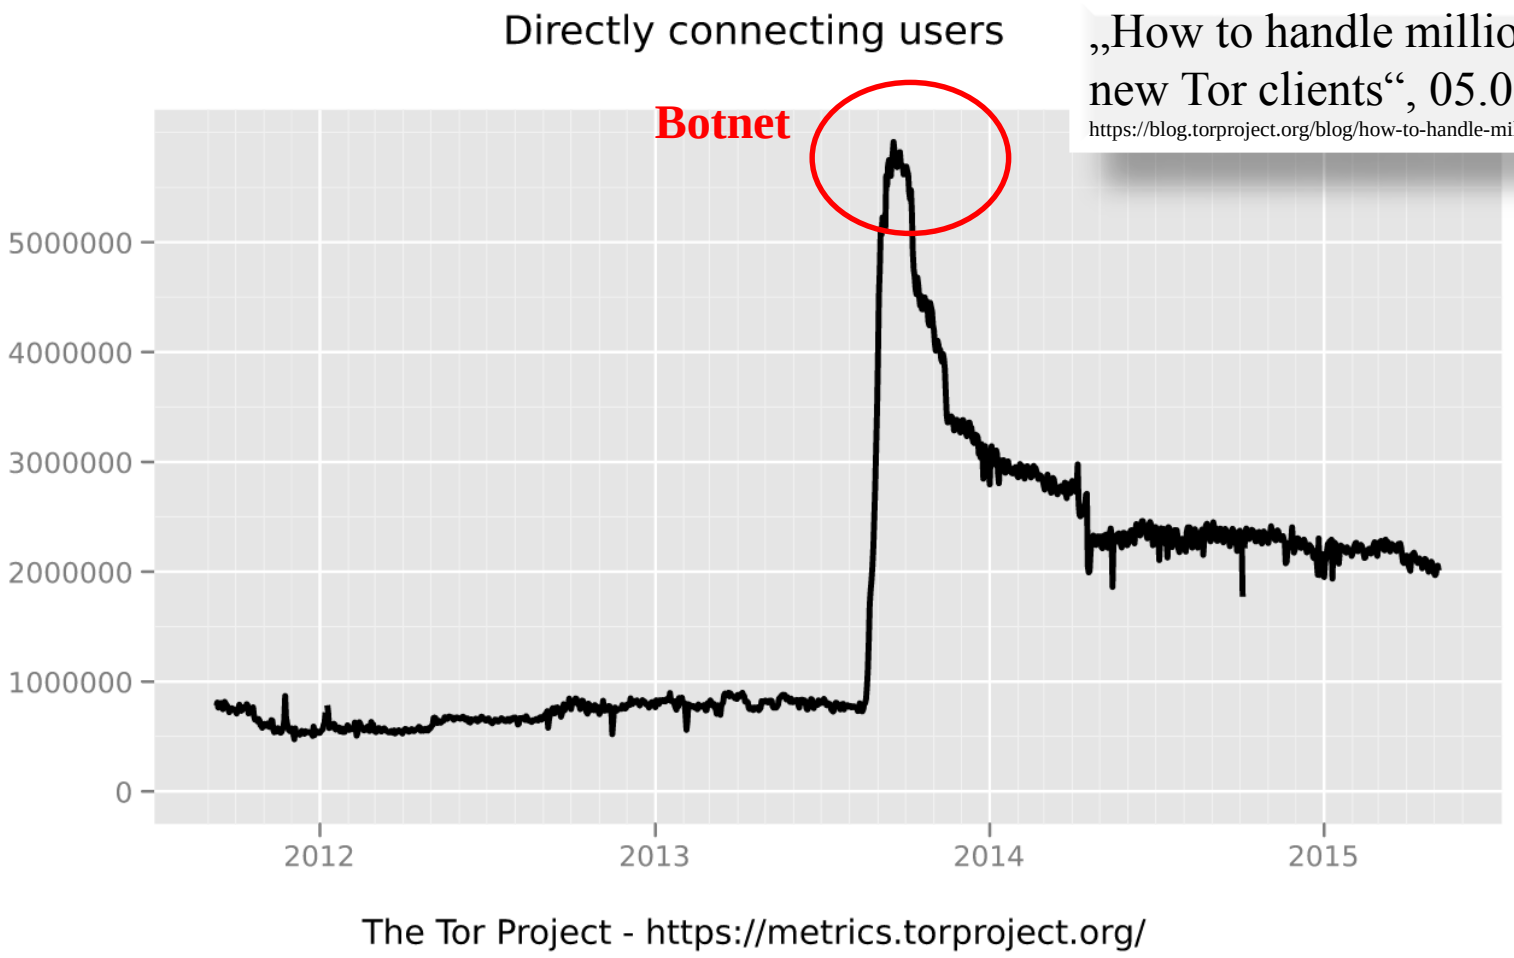

„How to handle millions of new Tor clients“, 05.09.2013
https://blog.torproject.org/blog/how-to-handle-millions-new-tor-clients

How it works... (1)

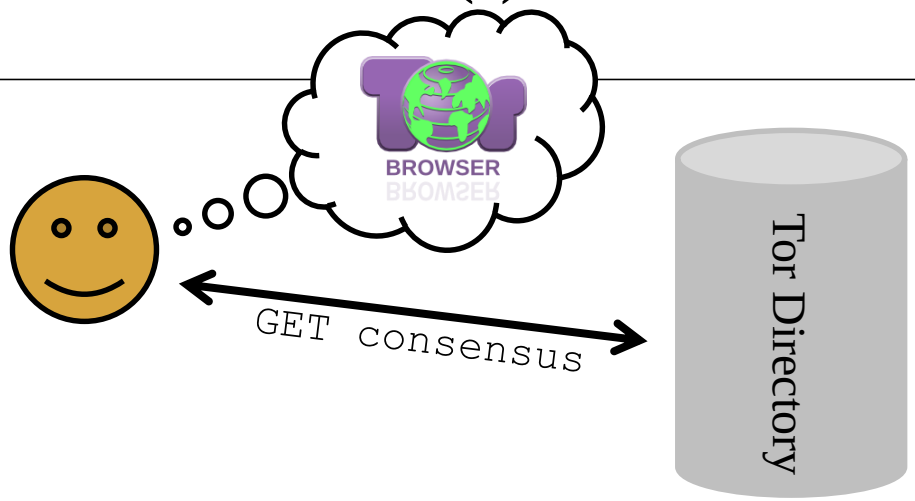

How it works... (2)

TECHNISCHE UNIVERSITÄT DARMSTADT

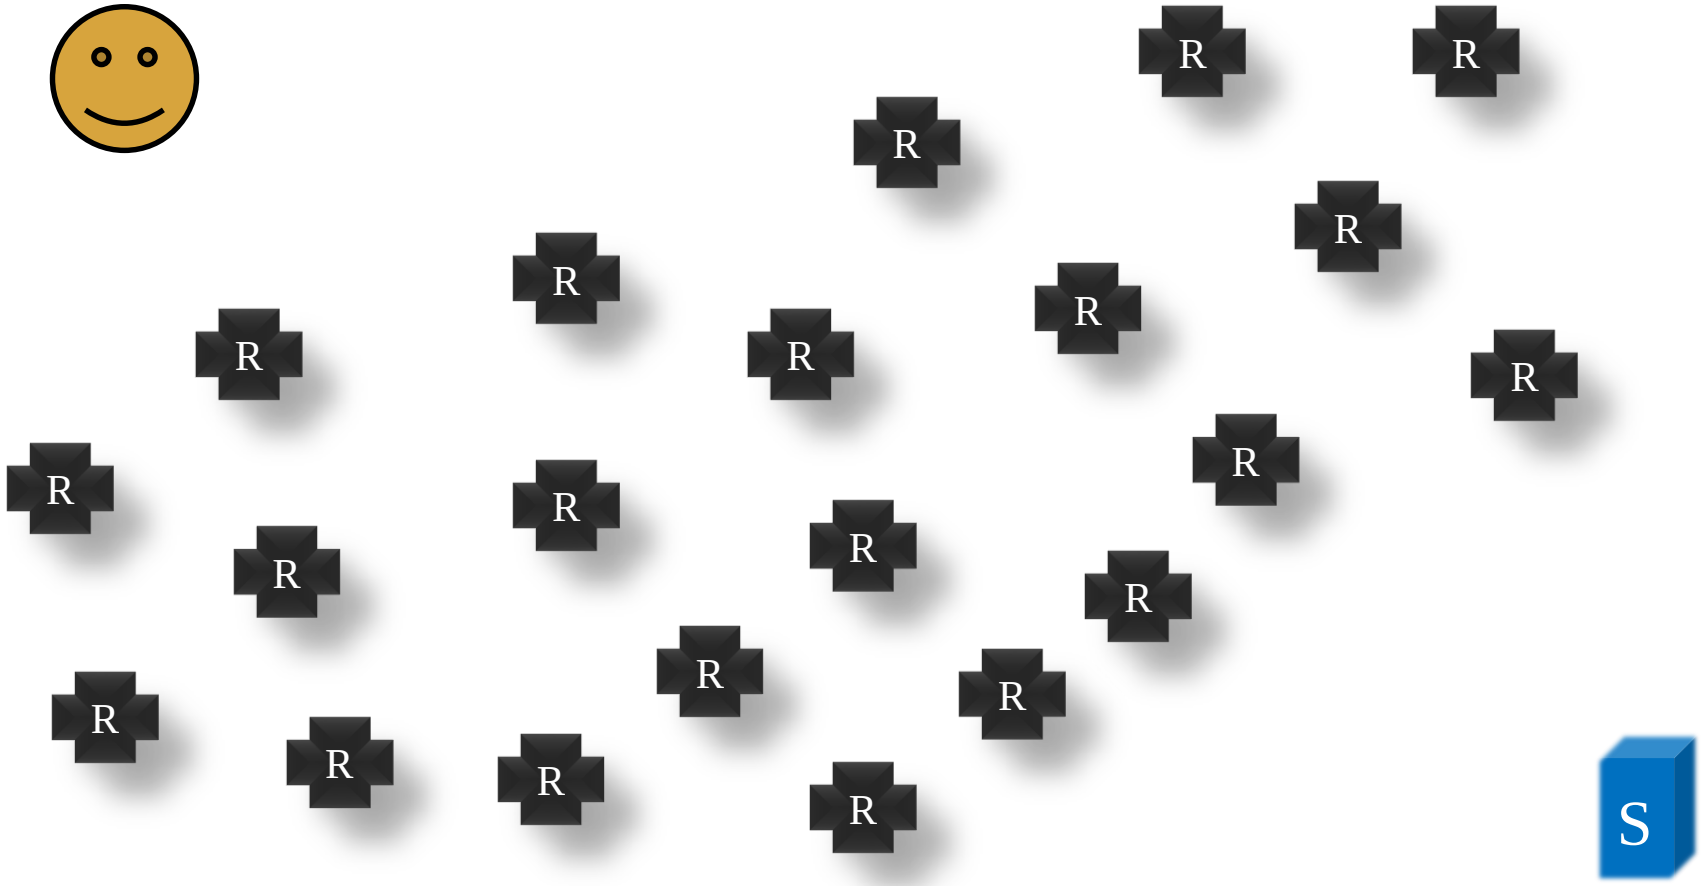

How it works... (3)

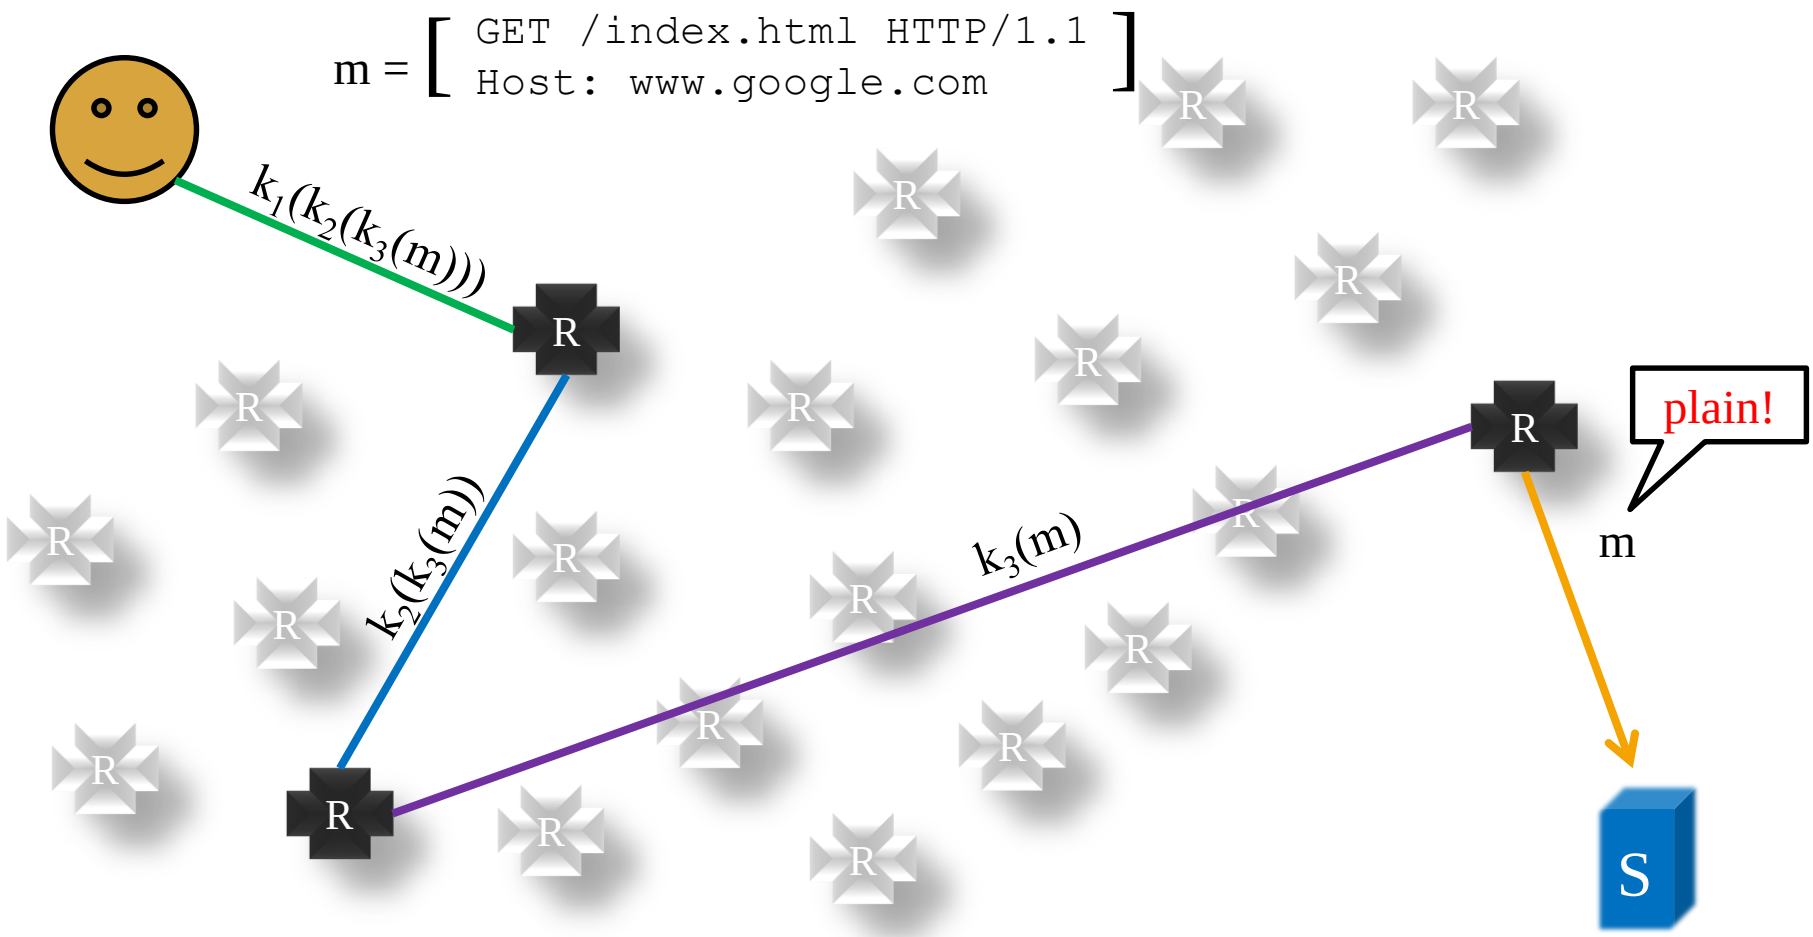

CollecTor
Consensuses and Server Descriptors

TECHNISCHE UNIVERSITÄT DARMSTADT

- Available at https://collector.torproject.org
- *Consensus* of the directory authorities
  - Published every hour
  - Defines network state as list of relays
- More details per relay in *Server Descriptors*
- Example entry of a consensus document:

```
r NotInMyBackyard 3B2fxLXY5M+0cu4Pvqgcv1cY7hY pBqKOtU+Wxk9GG6woIgoXZV0jU4 2015-05-01 16:47:18 87.106.21.77 9001 0
s Fast HSDir Running Stable Valid
v Tor 0.2.5.12
w Bandwidth=30
p reject 1-65535
```

Research
Privacy Engineering

TECHNISCHE UNIVERSITÄT DARMSTADT

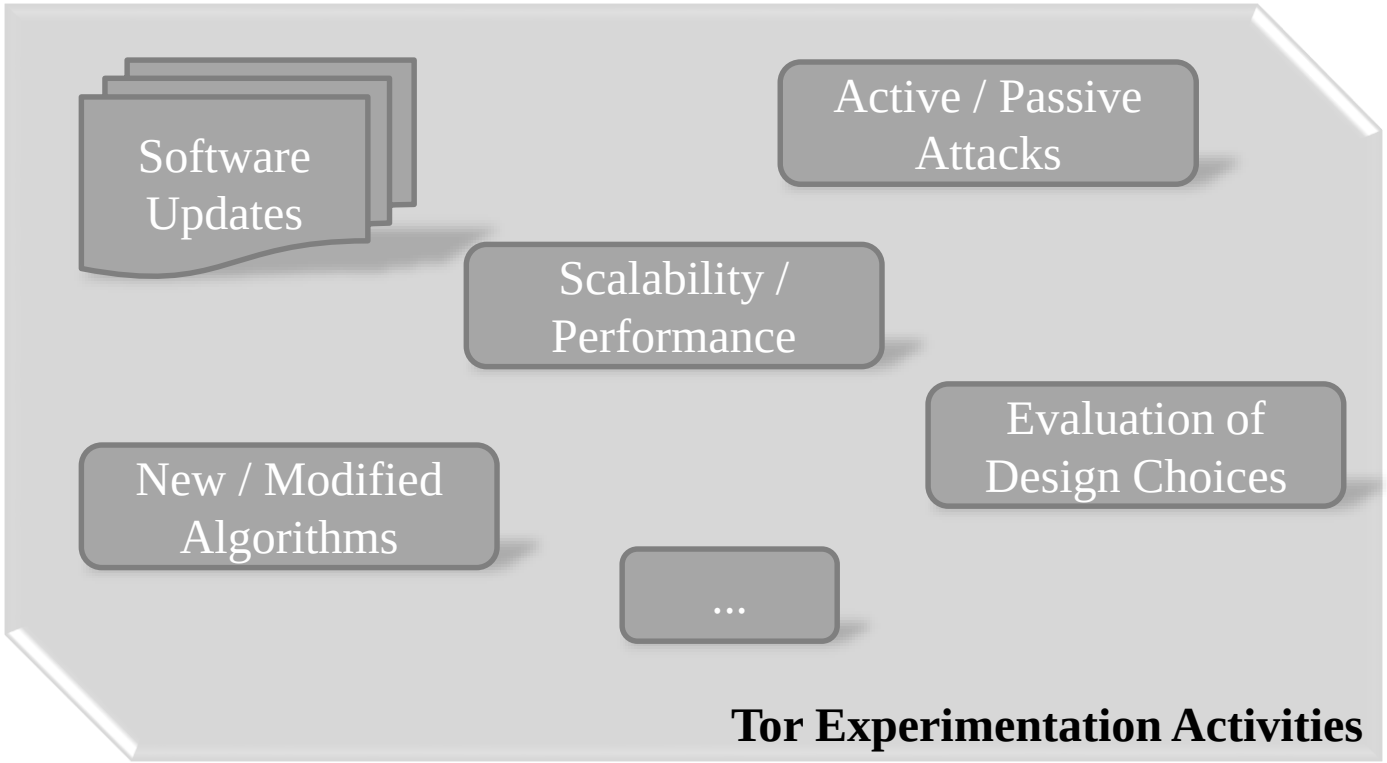

Experimentation is mandatory for privacy research on Tor!

TECHNISCHE UNIVERSITÄT DARMSTADT

Tor Experimentation Tools

- Live Experimentation
- Requirements
- Categorization
- Evaluation
- Simulation vs. Emulation

Live Experimentation

TECHNISCHE UNIVERSITÄT DARMSTADT

Advantages:

- Low costs
  - e.g. running a relay
- Easy to adapt / extend
  - Tor is open-source software
- Most realistic environment

Limitations:

- No control over the experiment
- Limited to deployed network
  - e.g. Tor software versions
- Results cannot be reproduced
- Might **threaten user's anonymity** and QoS [6]

Not Recommended (for most Experiments, see [6])

Safe & Realistic Environment Required

Requirements

TECHNISCHE UNIVERSITÄT DARMSTADT

Realism

Flexibility & Control

Safety

Scalability

Categorization – Evaluation

TECHNISCHE UNIVERSITÄT DARMSTADT

1. Live Tor Network
2. Analytical / Theoretical Modeling
3. Private Tor Networks
4. Overlay Testbed Deployments
5. Simulation
6. Emulation

Problematic!

Verification required

Limited scalability

Overlay Testbeds

TECHNISCHE UNIVERSITÄT DARMSTADT

- Services:
  - PlanetLab
  - Emulab
  - Deter
- Limitations:
  - Scalability
  - Results depend on current network state

    → cannot be reproduced (easily)
  - Shared resources

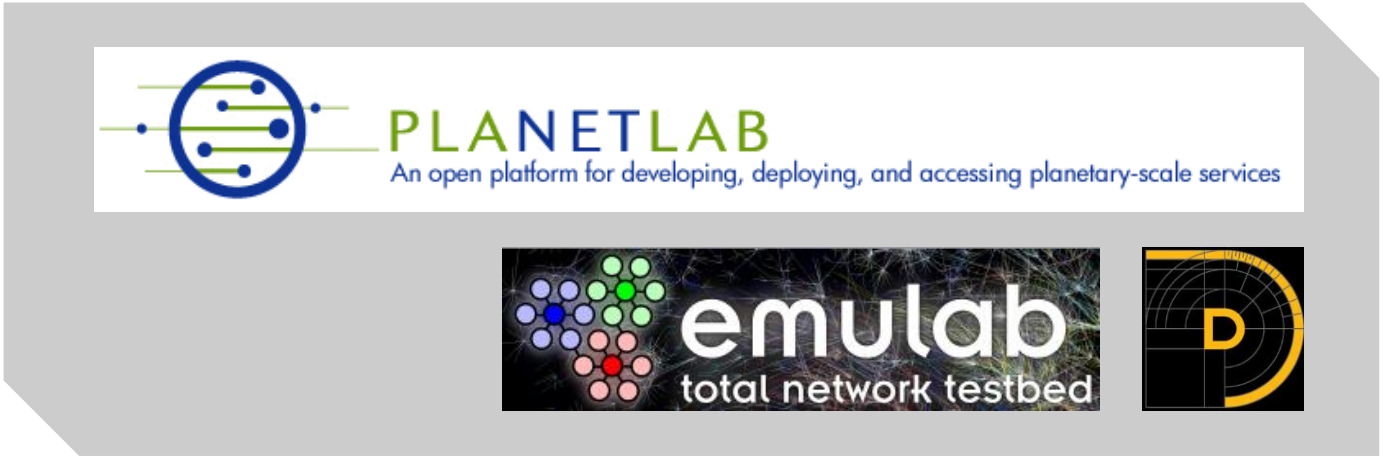

Categorization – Evaluation

TECHNISCHE UNIVERSITÄT DARMSTADT

1. Live Tor Network — Problematic!
2. Analytical / Theoretical Modeling — Verification required
3. Private Tor Networks — Limited scalability
4. Overlay Testbed Deployments — Results cannot be reproduced
5. Simulation
6. Emulation

Simulation vs. Emulation

TECHNISCHE UNIVERSITÄT DARMSTADT

Simulation

- Abstract model of the system, assumptions for simplicity
- Virtual time
- ➢ Reduced hardware requirements
- ➢ Improved scalability

Emulation

- Little to no assumptions, all operations performed
- Real time
- ➢ Substantial hardware requirements
- ➢ Scalability limited
  - Due to required hardware

Shadow | TorPS | COGS

ExperimenTor | SNEAC

TECHNISCHE UNIVERSITÄT DARMSTADT

Tor Experimentation Tools

- Metrics
- Simulators
  - Shadow, TorPS, COGS
- Emulators
  - ExperimenTor, SNEAC

Evaluation Metrics

TECHNISCHE UNIVERSITÄT DARMSTADT

1. Size / number of relays
2. Routing approach
3. Topology
4. Network effects (e.g. congestion)
5. Number of users
6. Usage patterns
7. Modeling adversaries

Experiment characteristics

8. Currently maintained?
9. Runs unmodified Tor source code?
10. Resource requirements

Tool characteristics

Shadow
[8] Jansen et al.

Simulator

TECHNISCHE UNIVERSITÄT DARMSTADT

- General-purpose, discrete-event simulator
- Runs on a single machine with user privileges
- Applications run as plugins
  - Tor plugin: Scallion

Limitations:

- Scalability limited by resources of a single host
- Simplifications might influence results, e.g.
  - Cryptographic operations are simulated by time delays
  - Downscaling of experiments

Shadow: Simulation Flow

Simulator

TECHNISCHE UNIVERSITÄT DARMSTADT

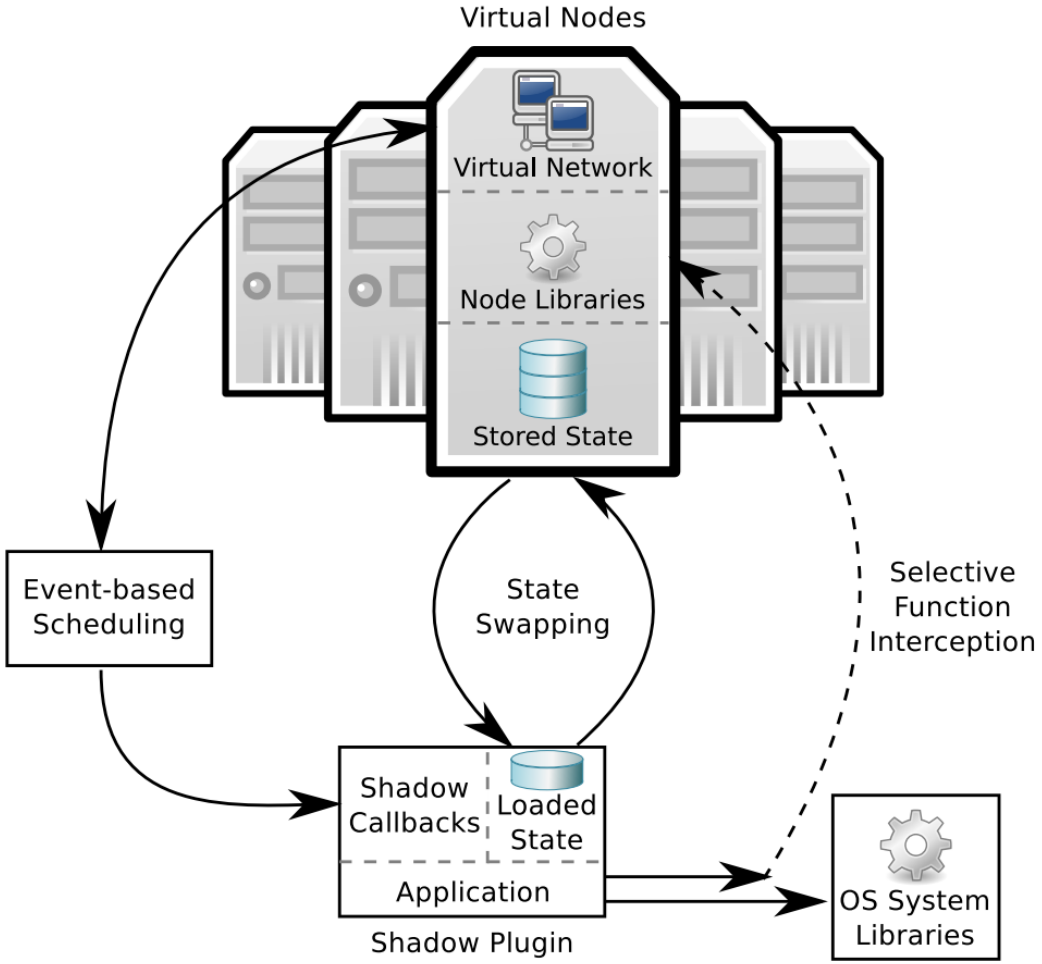

Source: [8]

Tor Path Simulator (TorPS)

[7] Johnson et al.

Simulator

TECHNISCHE UNIVERSITÄT DARMSTADT

- Specialized Tor simulator
- Simulate relay selection for circuit construction
- Intention: Test different algorithms

Limitations:

- Underlying network effects ignored
- Reimplementation of algorithms (python)

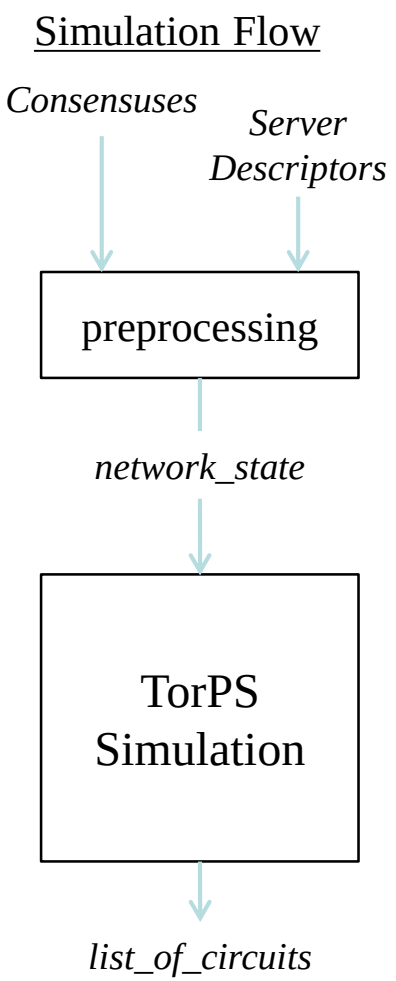

Changing of the Guards (COGS)

[5] Elahi et al.

Simulator

TECHNISCHE UNIVERSITÄT DARMSTADT

Purpose: Analyze effects of entry guard selection on user privacy

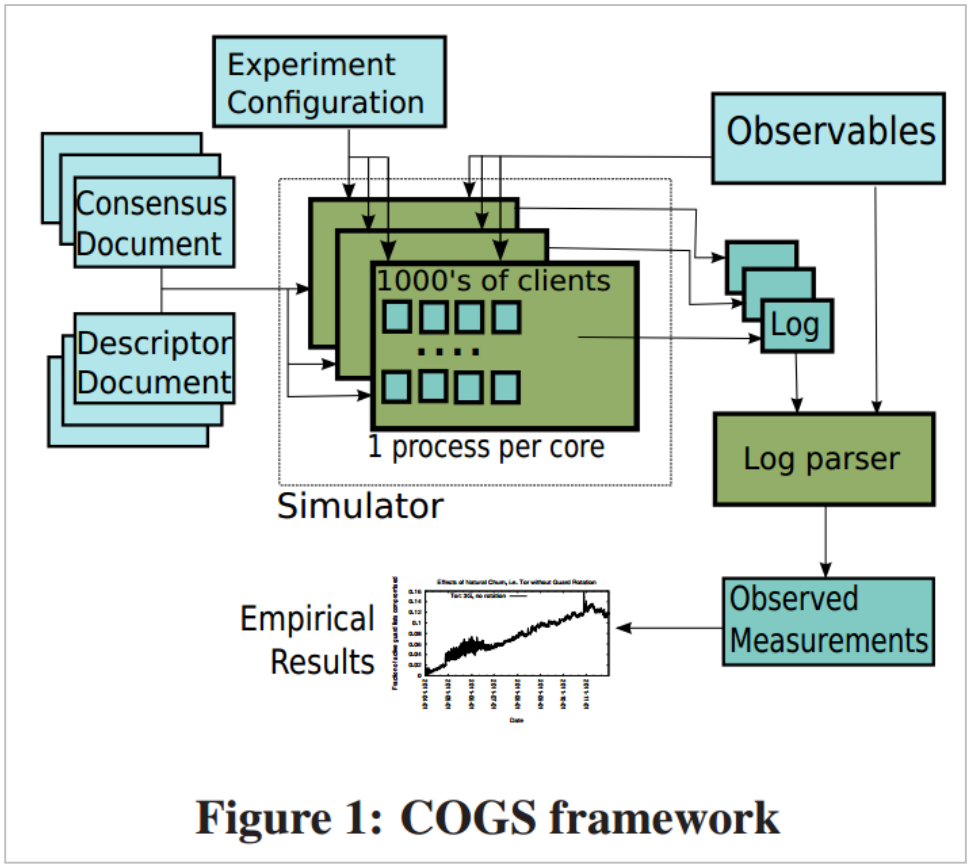

Figure 1: COGS framework

Source: [5]

ExperimenTor
[9] Bauer et al.

Emulator

TECHNISCHE UNIVERSITÄT DARMSTADT

- General-purpose Tor emulator
- At least two hosts required:
  - (Emulator core)+: Emulating the network topology
  - (Edge node)+: Running unmodified applications, e.g.
    - Web browsers, BitTorrent clients, ...

Limitations:

- Based on an outdated version of FreeBSD
- No longer available & maintained

→ Supposed to be replaced by SNEAC

ExperimenTor

Emulator

TECHNISCHE UNIVERSITÄT DARMSTADT

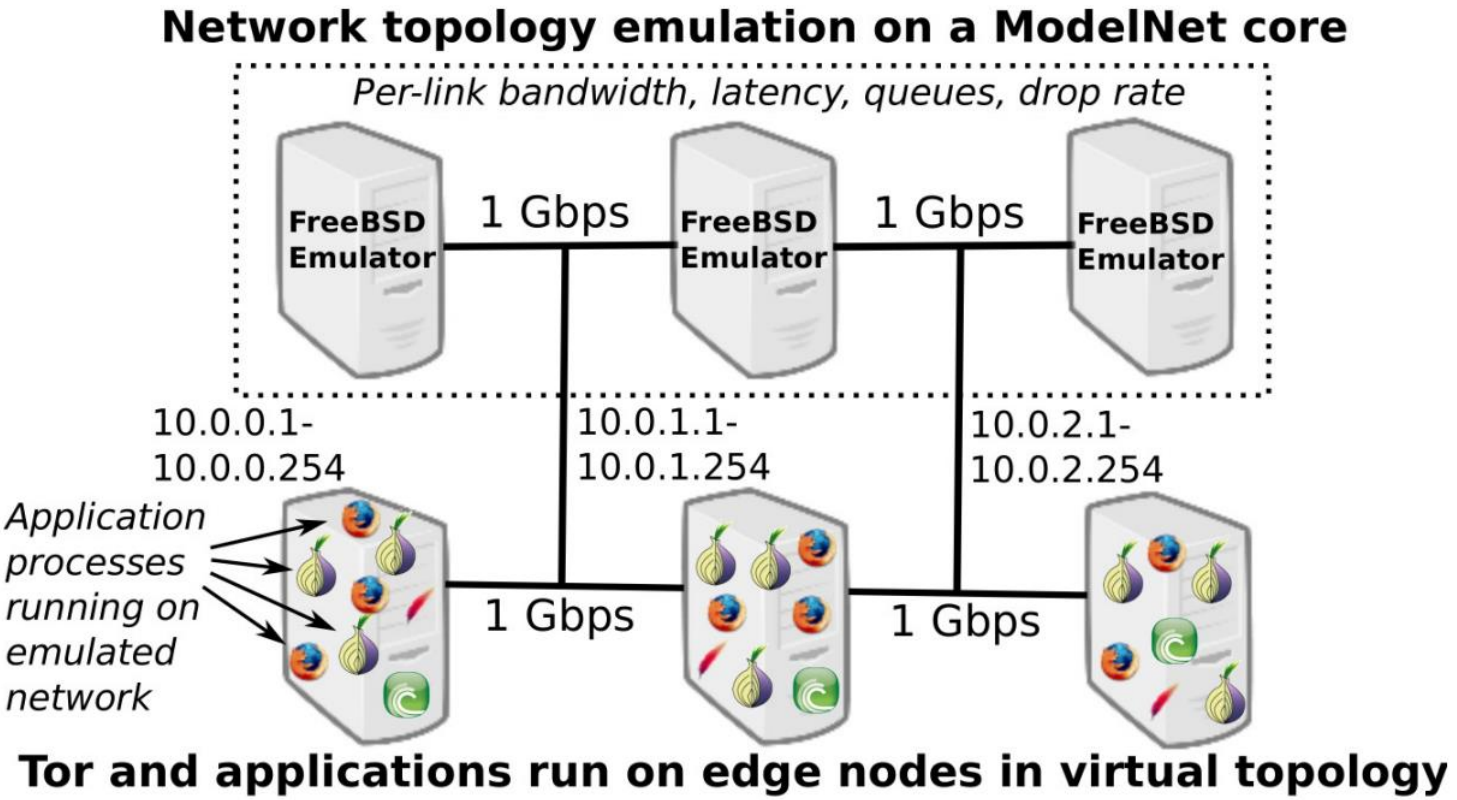

Figure 2: ExperimenTor system architecture

Source: [9]

SNEAC
[18] Singh

Emulator

TECHNISCHE UNIVERSITÄT DARMSTADT

Scalable Network Emulator for Anonymous Communication

Limitations:

- Hardware requirements lim
  scalability!
- Requires own data extractic
- User Model?

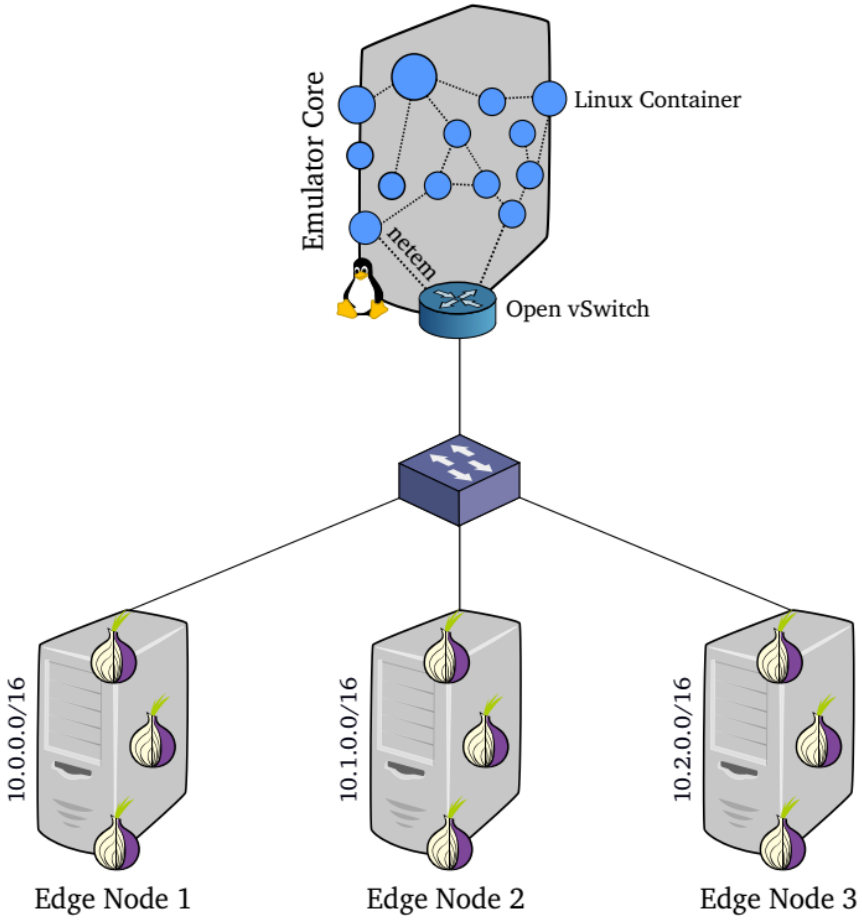

Figure 3.1: Architecture of a SNEAC setup being used to emulate Tor.

Source: [18]

Comparison

TECHNISCHE UNIVERSITÄT DARMSTADT

Metric	Shadow	TorPS	ExperimenTor
1. Size / number of relays	downscaling, simulation with 500+ relays possible	no downscaling	limited by available resources
2. Routing approach	not using additional weighting in node selection	ignoring paths being dropped due to timeouts	-
3. Topology	geographic distribution ignored, bandwidth distribution based on Tor	both same as Tor	geographic distribution of Tor ignored, bandwidth distribution based on Tor
4. Network effects (e.g. congestion)	yes	no	yes (simplified)
5. Number of Tor users	downscaled	no	downscaled
6. Usage pattern of Tor users	5 usage patterns	5 usage patterns	2 usage patterns
7. Modeling adversaries	possible	possible	possible
8. Currently being maintained	yes	yes	no
9. Using original Tor code	yes	no, Python application	yes
10. Required resources	single host, user privileges	single host, user privileges	min. 2 hosts, high resource requirements

Conclusion

TECHNISCHE UNIVERSITÄT DARMSTADT

- No standardized experimentation approach
  - Simulation vs. emulation
- Experimentation results are based on specific tools

  → cannot be compared easily
- Inherent complications experimenting with an anonymity network
- General problems:
  - User model / traffic
  - Scalability / downscaling

Thank you for your attention!

Questions?

Matthias Göhring
de.m.goehring@ieee.org

Acknowledgements

The authors would like to thank

- Rob Jansen
- Aaron Johnson
- Ian Goldberg
- Kevin Bauer
- Sukhbir Singh

References

TECHNISCHE UNIVERSITÄT DARMSTADT

[5] T. Elahi, K. Bauer, M. AlSabah, R. Dingledine, and I. Goldberg, "Changing of the guards: A framework for understanding and improving entry guard selection in tor," in Proceedings of the Workshop on Privacy in the Electronic Society (WPES 2012), ACM, October 2012.

[6] K. Loesing, S. J. Murdoch, and R. Dingledine, "A case study on measuring statistical data in the Tor anonymity network," in Proceedings of the Workshop on Ethics in Computer Security Research (WECSR 2010), LNCS, Springer, January 2010.

[7] A. Johnson, C. Wacek, R. Jansen, M. Sherr, and P. Syverson, "Users get routed: Traffic correlation on tor by realistic adversaries," in Proceedings of the 2013 ACM SIGSAC Conference on Computer and Communications Security, CCS '13, ACM, 2013.

[8] R. Jansen and N. Hopper, "Shadow: Running tor in a box for accurate and efficient experimentation.," in Proceedings of the Network and Distributed System Security Symposium - NDSS'12, The Internet Society, 2012.

[9] K. Bauer, D. Mccoy, M. Sherr, and D. Grunwald, "Experimentor: A testbed for safe and realistic tor experimentation," in In: Proceedings of the USENIX Workshop on Cyber Security Experimentation and Test (CSET), 2011.

[18] S. Singh, "Large-scale emulation of anonymous communication networks," Master's thesis, University of Waterloo, 2014.