U.S. NAVAL RESEARCH LABORATORY

Inside Job: Applying Traffic Analysis to Measure Tor from Within

*Rob Jansen, *U.S. Naval Research Laboratory*
*Marc Juarez, *imec-COSIC KU Leuven*
Rafael Gálvez, *imec-COSIC KU Leuven*
Tariq Elahi, *imec-COSIC KU Leuven*
Claudia Diaz, *imec-COSIC KU Leuven*

**equally credited authors*

Rob Jansen
Center for High Assurance Computer Systems
U.S. Naval Research Laboratory

25th Symposium on Network and Distributed System Security
San Diego, CA
February 21st, 2018

Tor Website Fingerprinting

- **Adversary's goal: use website fingerprinting to deanonymize client (link client to destination)**

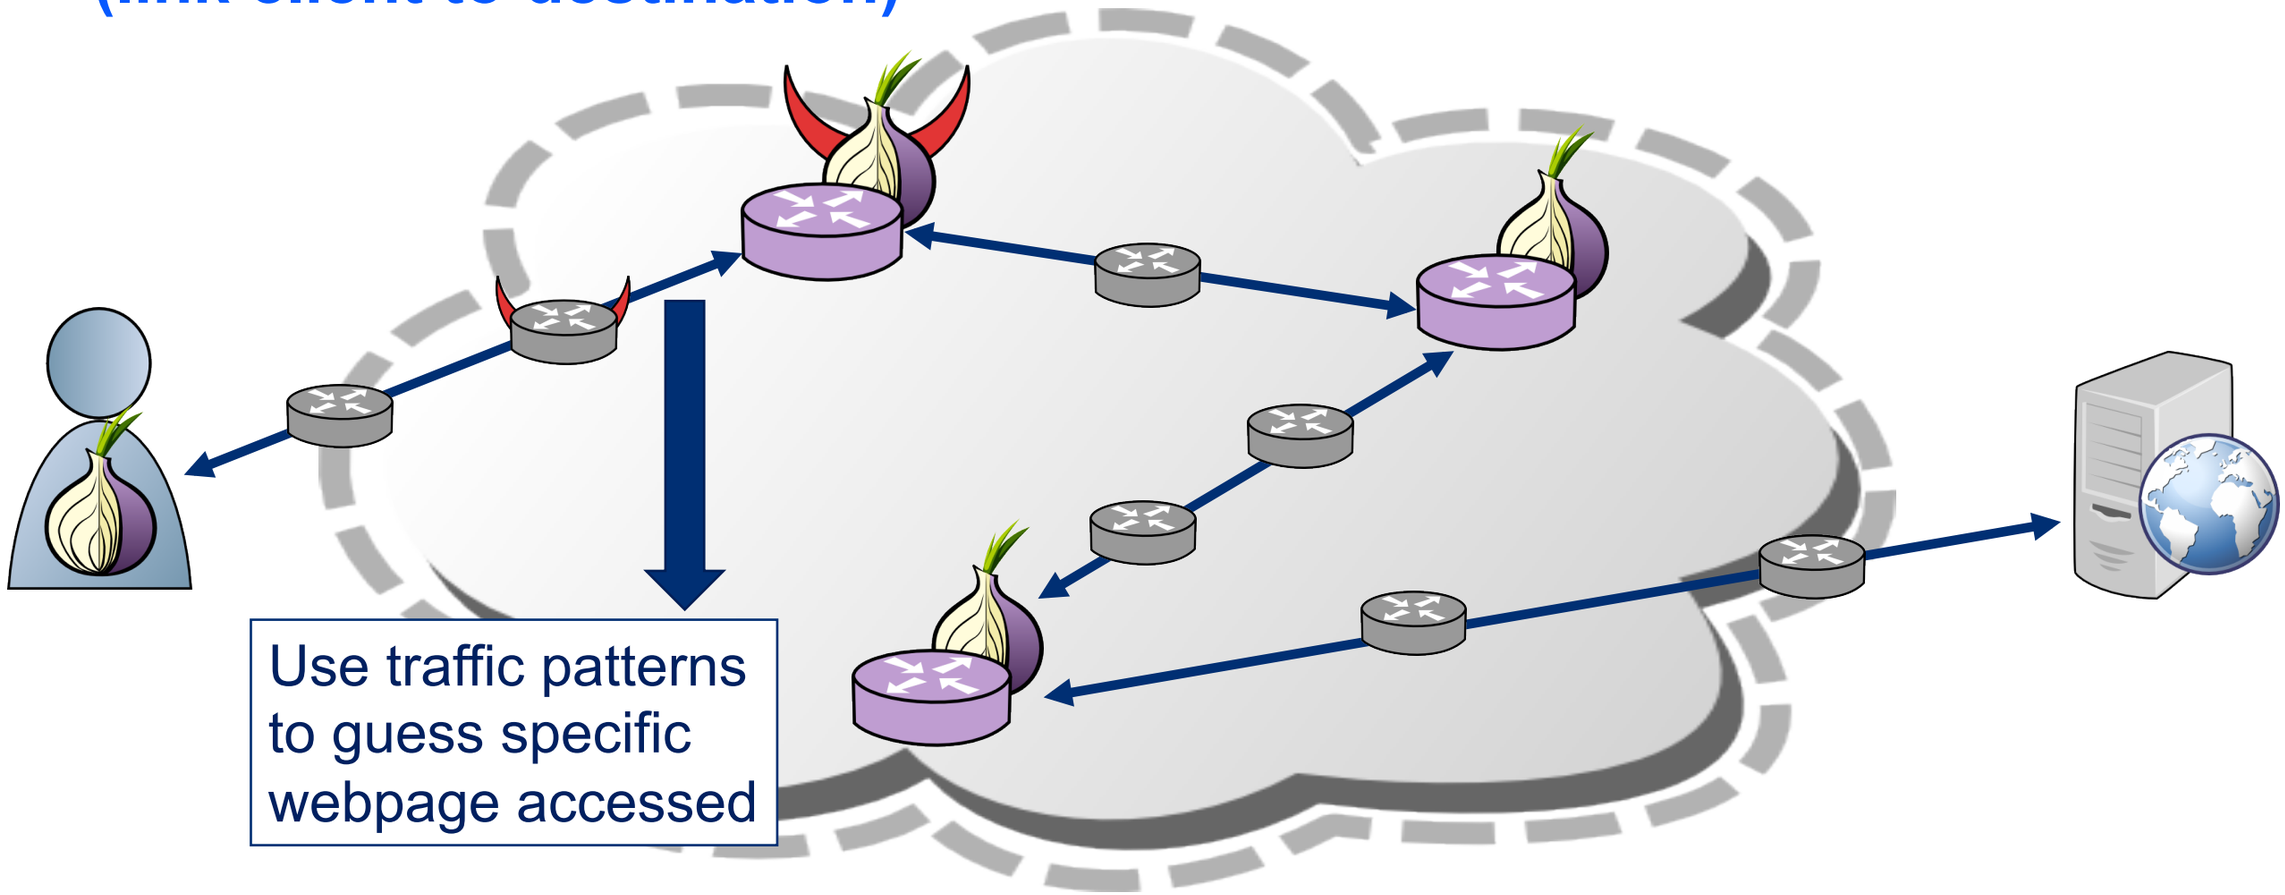

Onion Service Fingerprinting

- **Tor website fingerprinting on onion services**

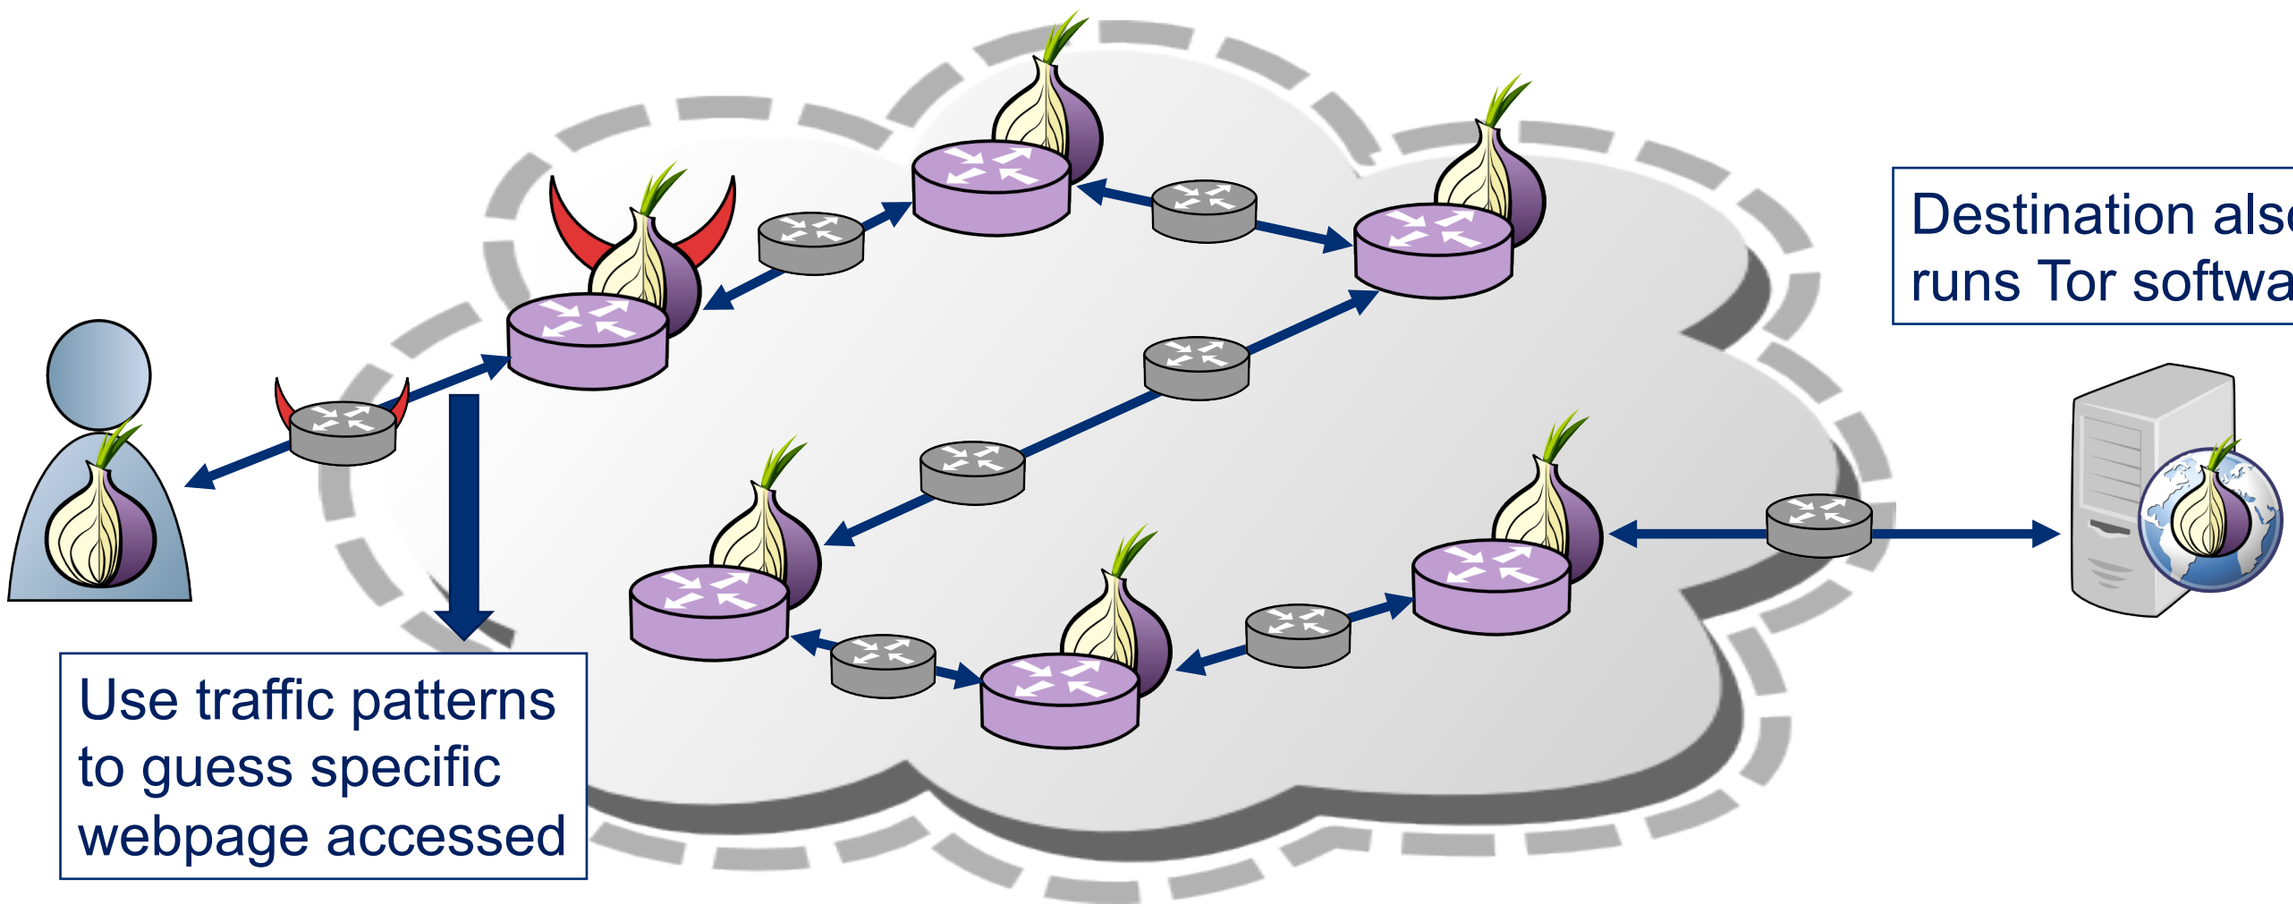

Onion Service Fingerprinting

- All prior work considers adversary in an entry position

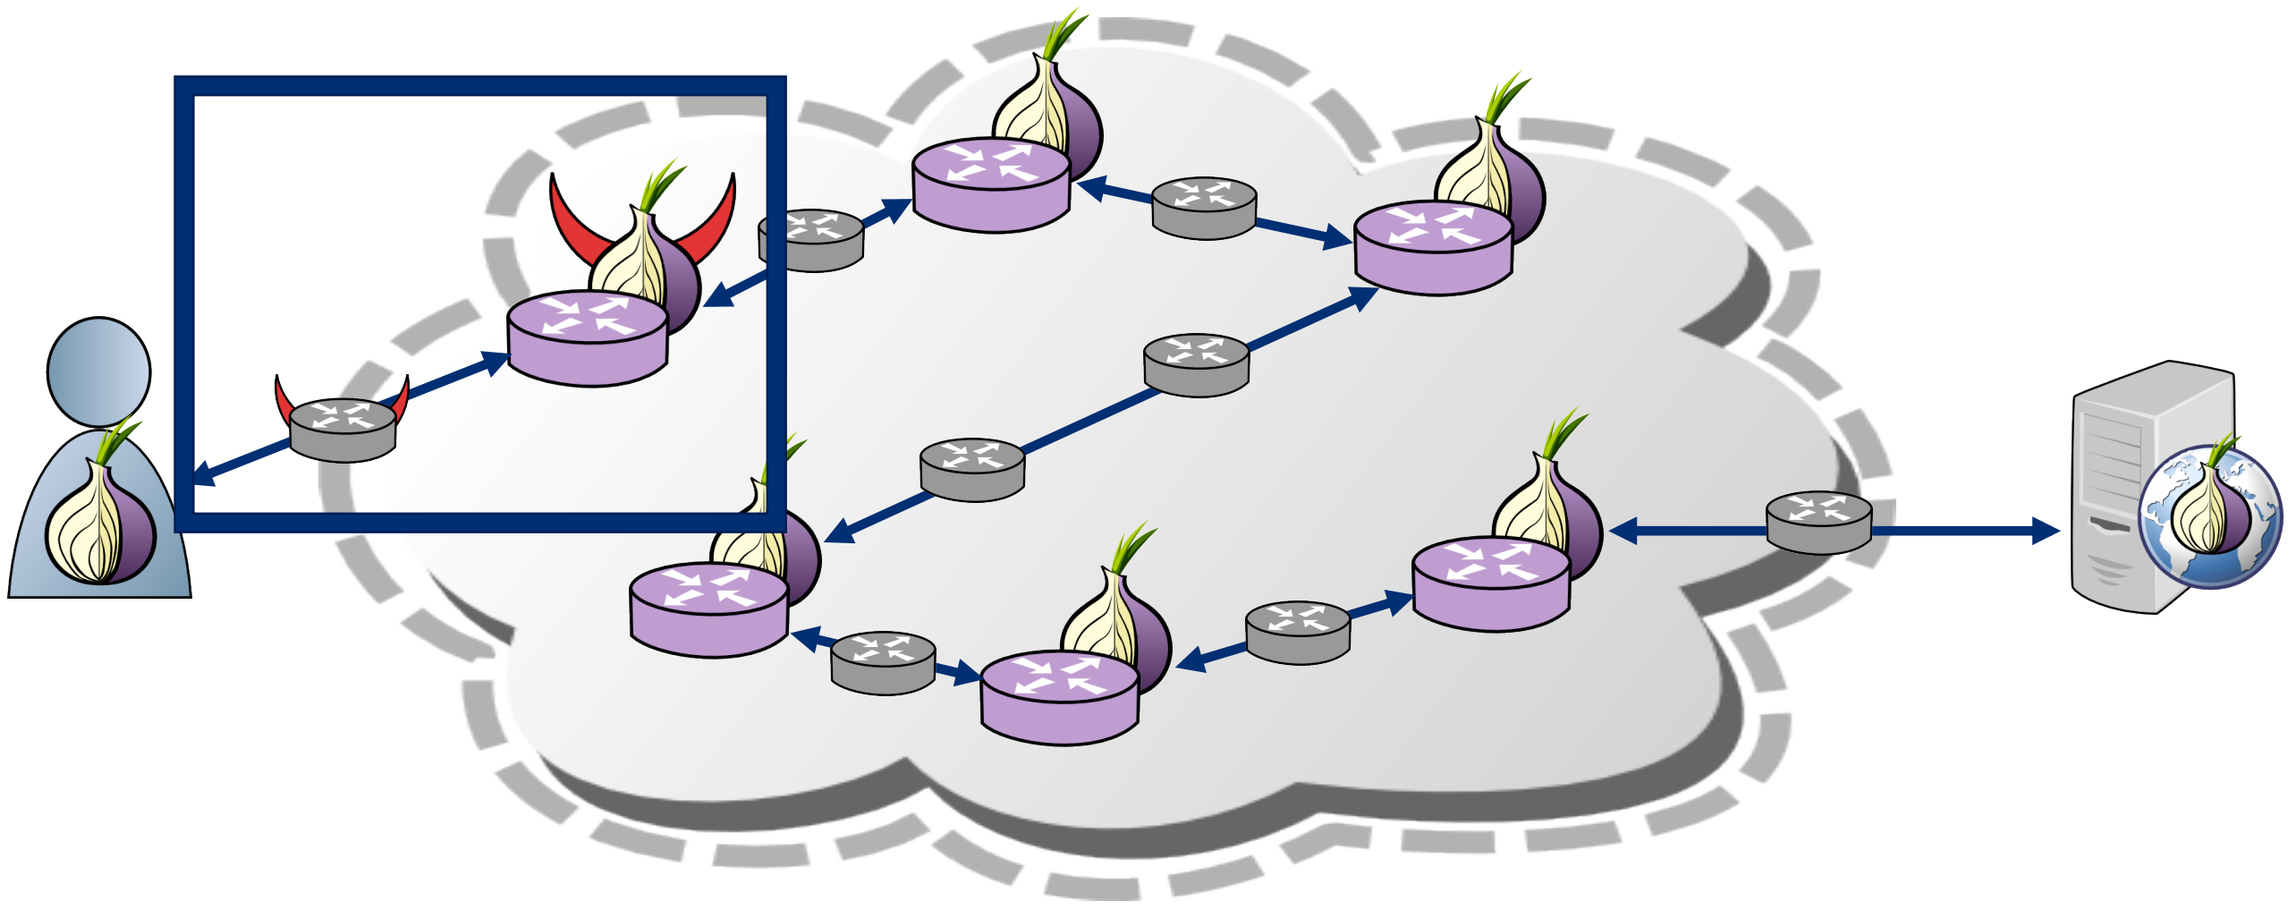

Onion Service Fingerprinting

- **All prior work considers adversary in an entry position**

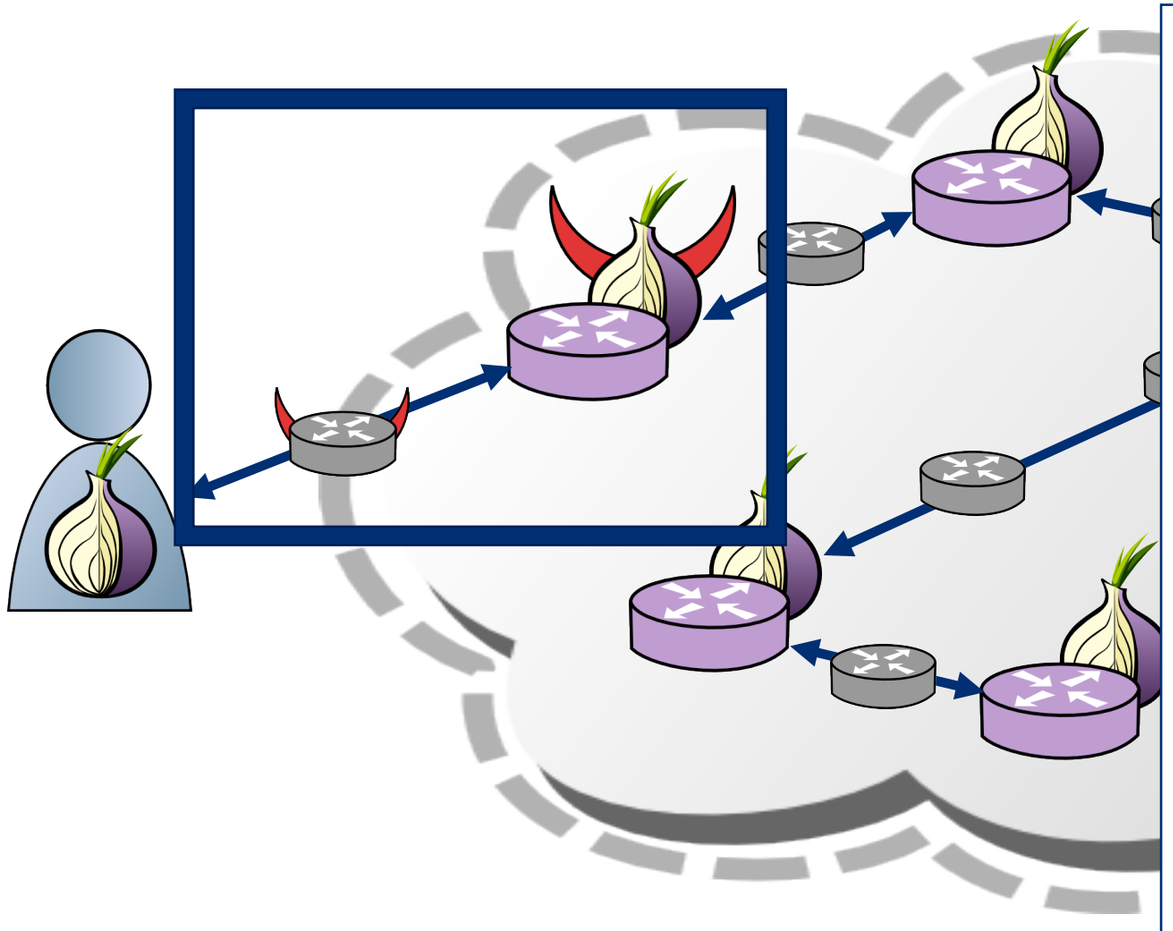

Limitations of the entry

- Client-to-entry path is an unrealistic privileged position for most
- Entry guard relays must be stable and have high up-time
- Clients choose and pin 1 entry guard for 2-3 months before switching
- It takes entry guards 3 months to reach steady state and be fully utilized by the network

This work: fingerprinting from middle relays

- **Onion service fingerprinting from an internal, middle relay position**

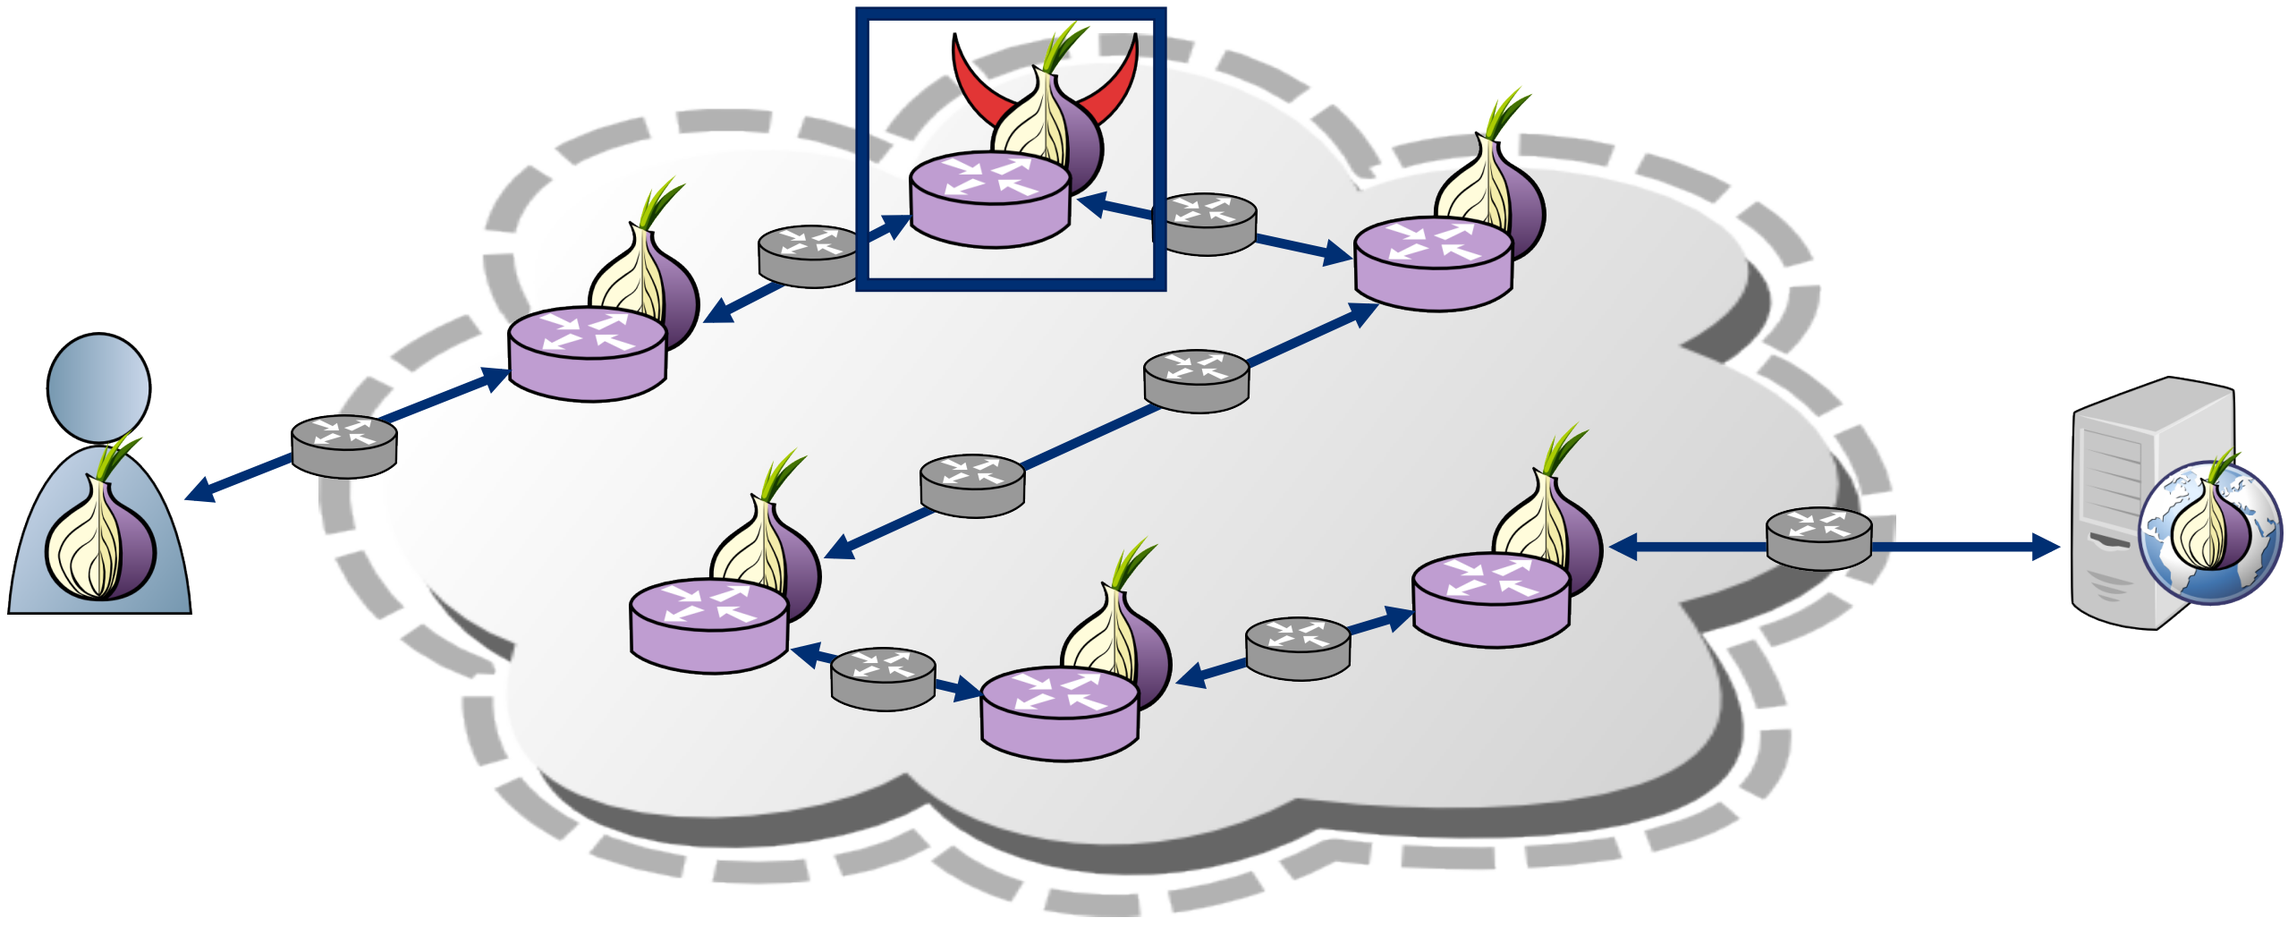

This work: fingerprinting from middle relays

- **Onion service fingerprinting from an internal, middle relay position**

Advantages of the middle

- Clients choose a new middle for every circuit (choice is weighted by bandwidth)
- No special relay requirements
- Fully utilized almost immediately
- Statistical sampling of all clients

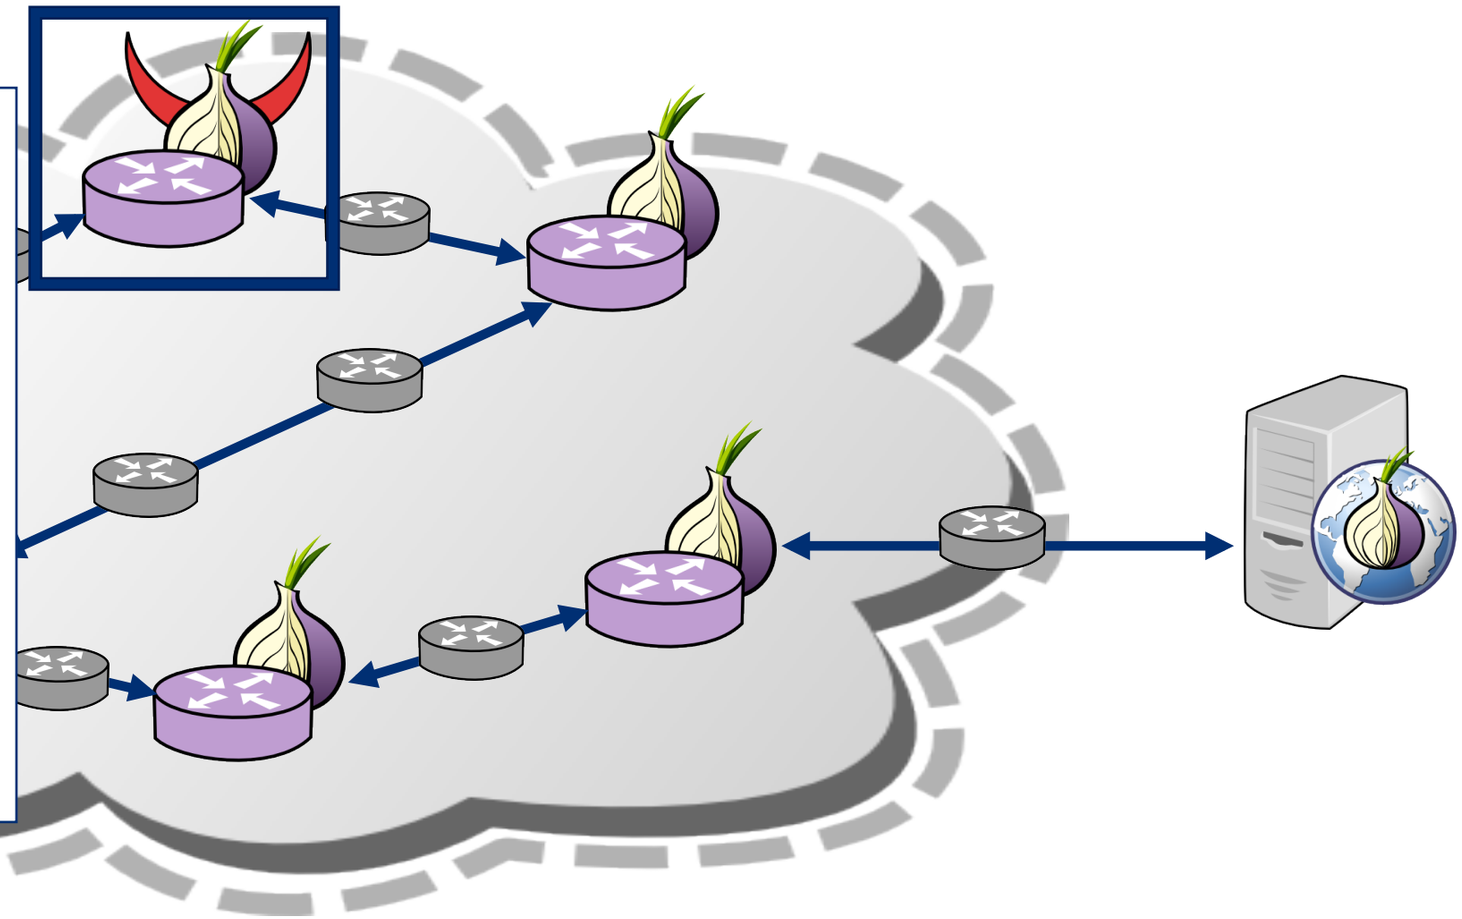

This work: fingerprinting from middle relays

- **Onion service fingerprinting from an internal, middle relay position**

Advantages of the middle

- Clients choose a new middle for every circuit (choice is weighted by bandwidth)
- No special relay requirements
- Fully utilized almost immediately
- Statistical sampling of all clients

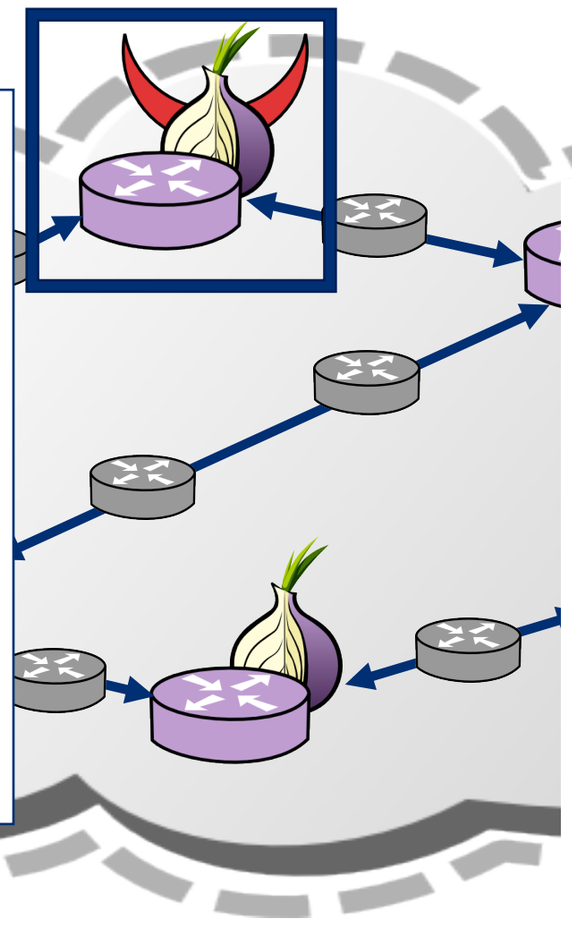

Middles will observe a client two orders of magnitude more quickly than guards

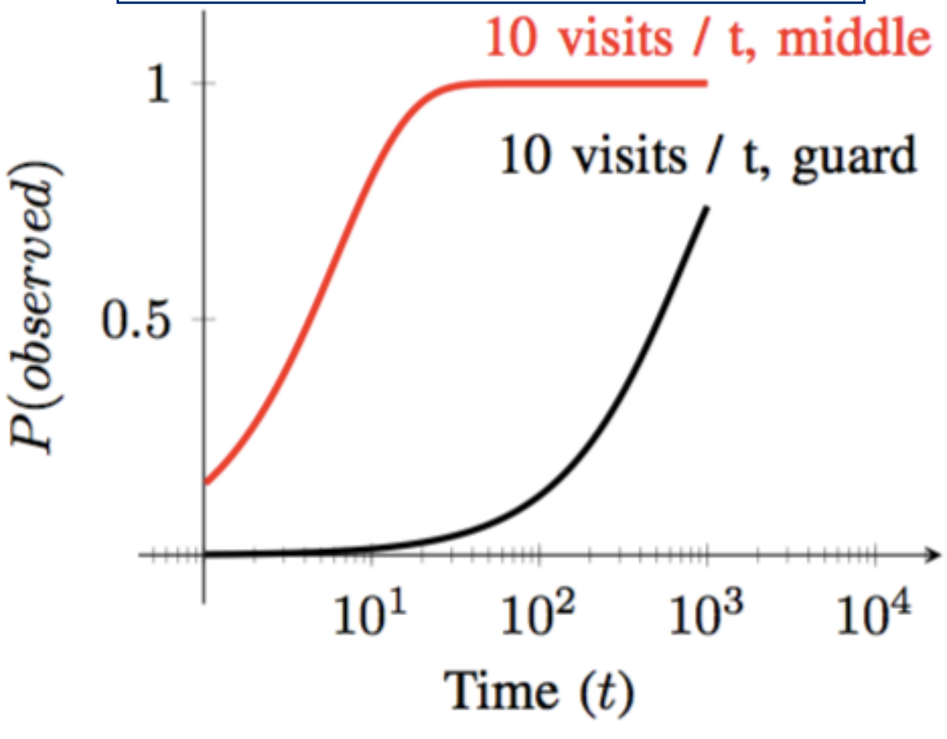

This work: fingerprinting from middle relays

- **The middle identifies the destination… and then what?**

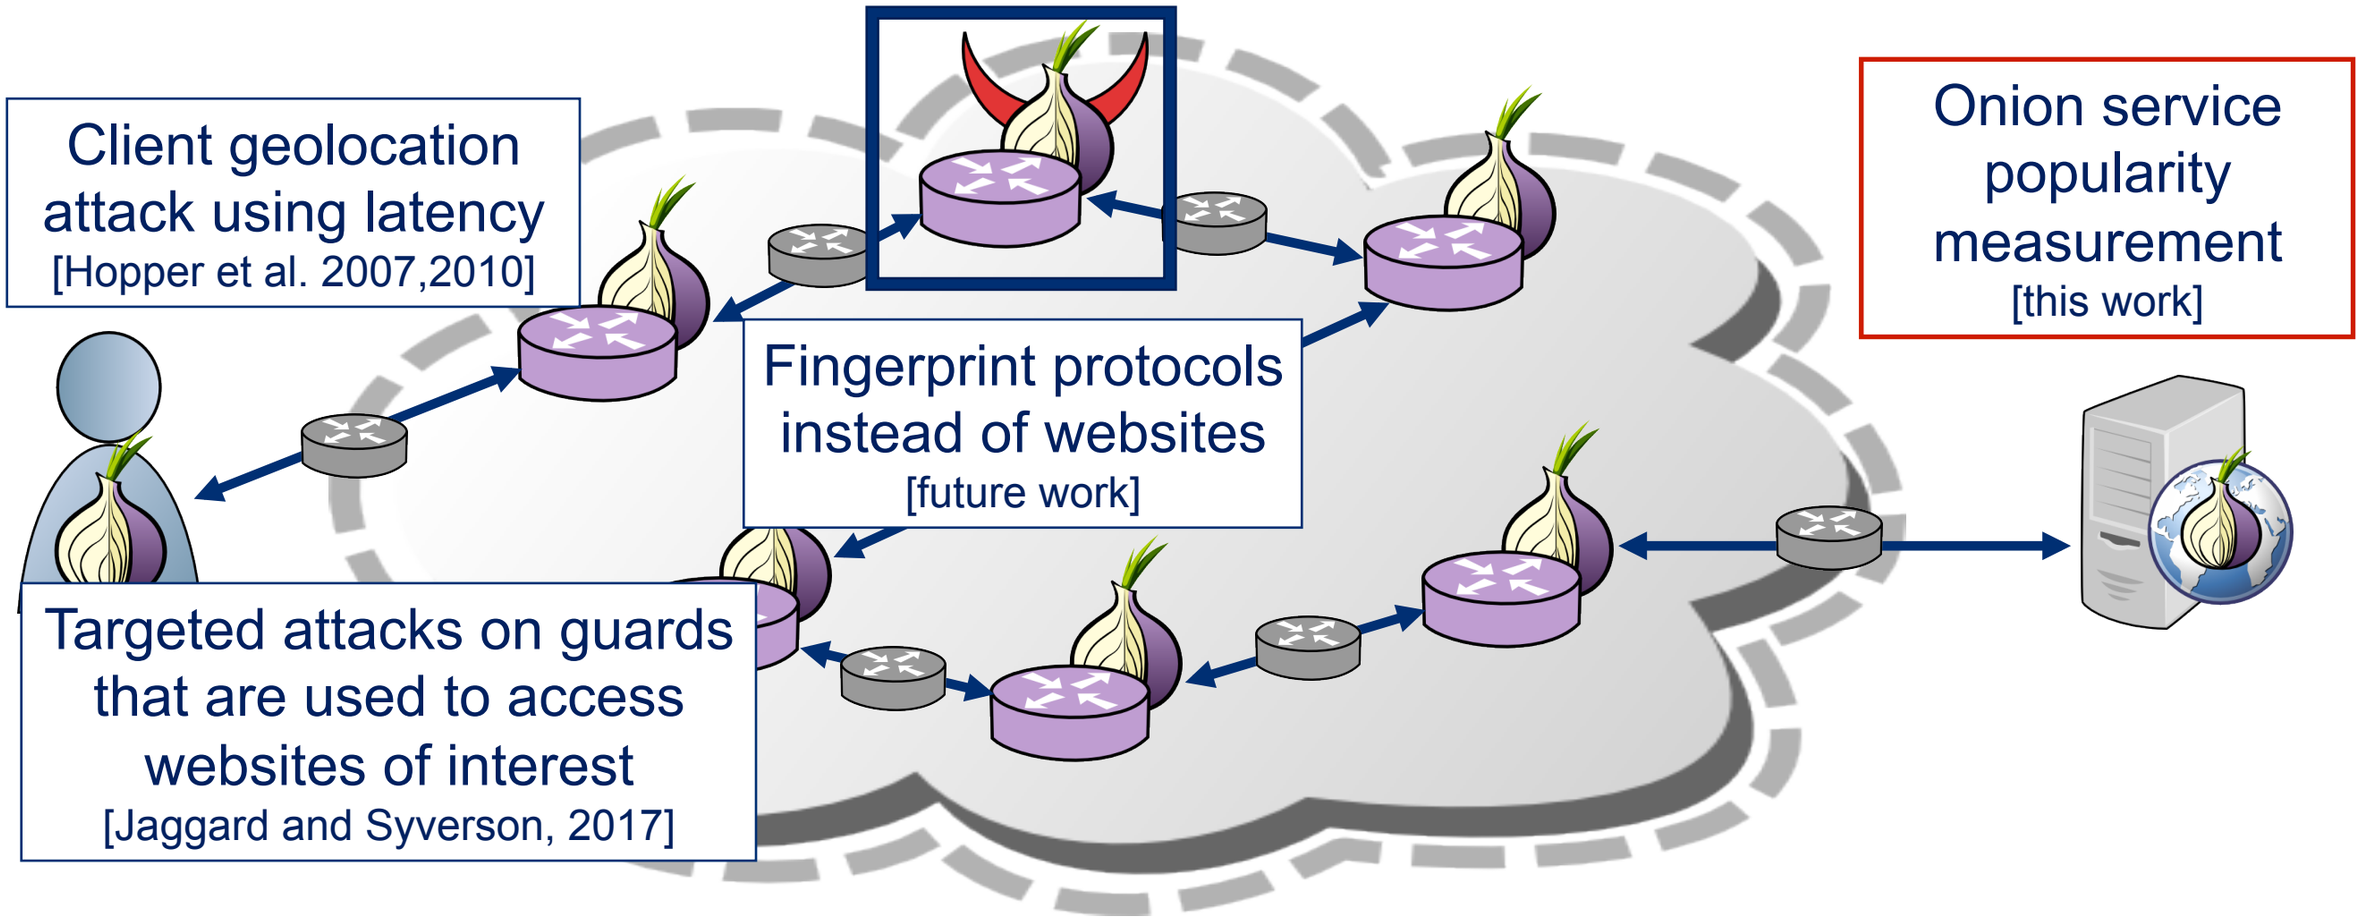

U.S. NAVAL RESEARCH LABORATORY

Outline

- ~~Background, Motivation: Why the middle relay?~~
- Circuit fingerprinting
- Onion Service Fingerprinting
- Onion Service Popularity Measurement
- Conclusion / Questions

Circuit Fingerprinting

- Collect circuit traces, extract features, train classifiers
- Identify circuit purpose and position

Circuit Fingerprinting

- **Predict circuit type and relay position**

Binary classification

- Circuit purpose:
  rendezvous (onion service)
- Circuit position:
  middle (adj. to guard)

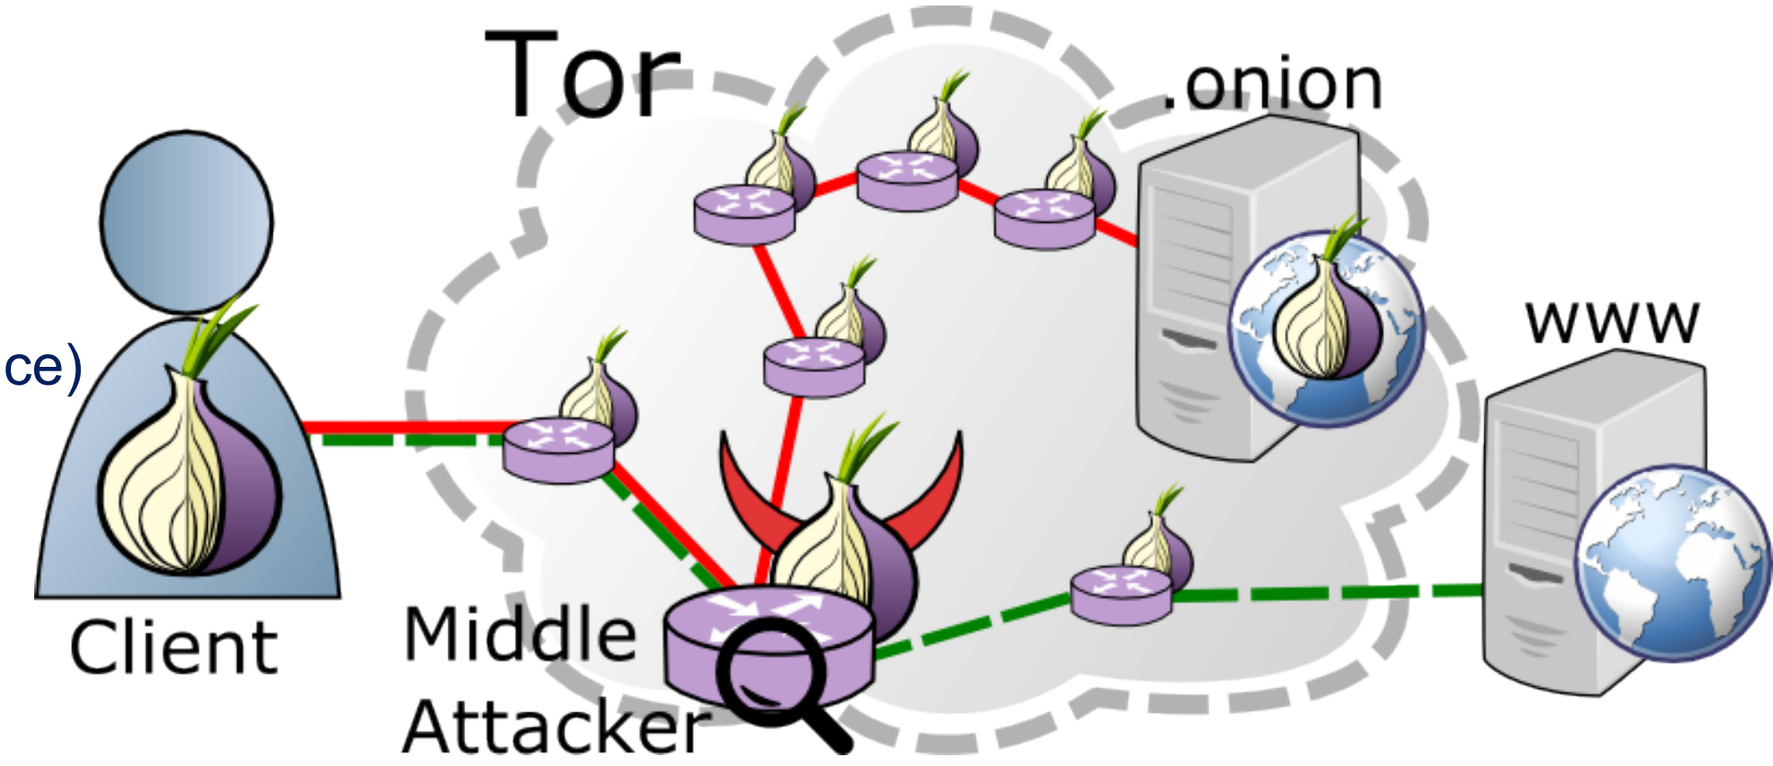

Data Set, Features, and Training

- **Generate samples using Shadow**
  - Use the Shadow Tor simulator to generate 1.85 million circuits
  - Label circuits with purpose and position
  - Extract features and train random-forest classifiers
- **Use as features:**
  - Previous/next node type
  - Counts of cell type/relay command (recv/sent inside/outside)

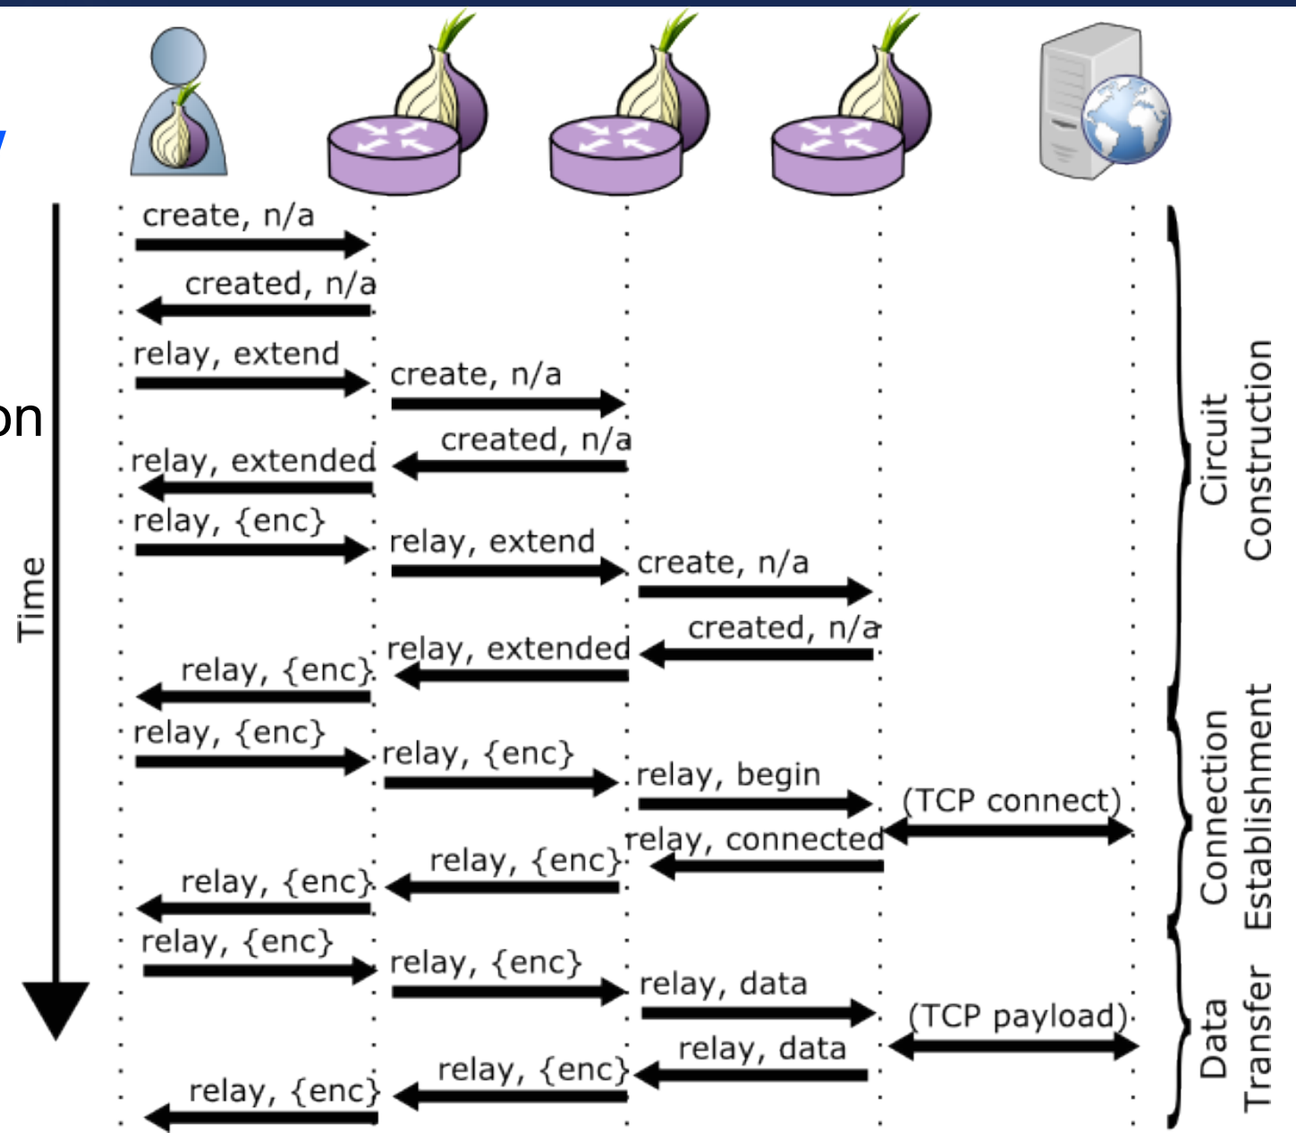

Circuit fingerprinting results

TABLE I. 10-FOLD CROSS-VALIDATED CIRCUIT CLASSIFICATION RESULTS

	Purpose (rendezvous vs other)	Position (C-M1 vs other)
Accuracy	$92.41 \pm 0.07\%$	$98.48 \pm 0.01\%$
Precision	$91.87 \pm 0.11\%$	$97.16 \pm 0.03\%$
Recall	$93.05 \pm 0.09\%$	$99.88 \pm 0.01\%$
F-1	$92.46 \pm 0.07\%$	$98.50 \pm 0.01\%$
True Positives	396,615 (91.77%)	821,478 (97.08%)
False Positives	35,576 (8.23%)	24,689 (2.92%)
False Negatives	30,056 (6.95%)	984 (0.12%)
True Negatives	402,135 (96.05%)	845,183 (99.88%)

Onion Service Fingerprinting

- Collect webpage traces, train and evaluate classifiers
- Identify onion service

Onion Service Fingerprinting

- **Given a rendezvous circuit, can we identify the destination?**

- Run modified Tor to identify **circuits that we originate!**
- Crawl known and online onion sites

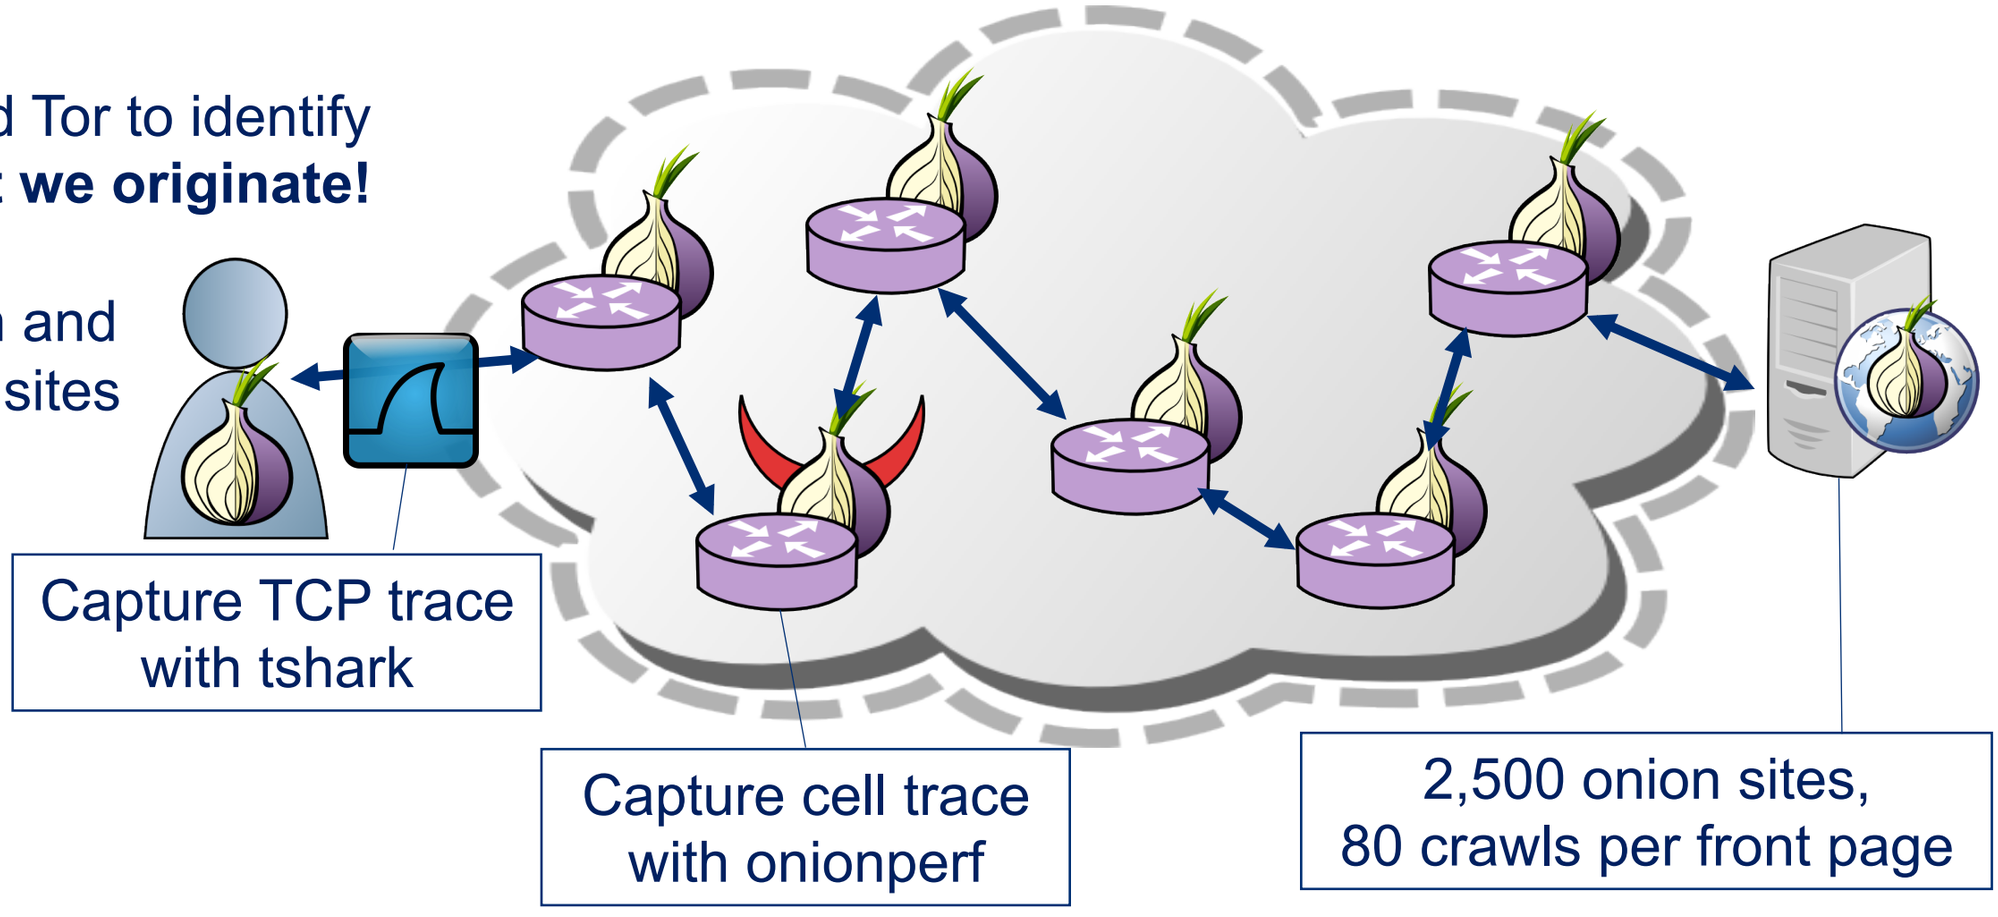

Closed World Onion Site Fingerprinting Results

- **True Positive Rates**

Entry model
- Classify using client-to-guard packet traces

Num sites	k-NN (%)	k-FP (%)	CUMUL (%)
10	95% ± 0.03	95% ± 0.06	92% ± 0.04
50	75% ± 0.02	85% ± 0.03	81% ± 0.02
100	67% ± 0.01	68% ± 0.03	64% ± 0.02

Middle relay model
- Classify using middle relay cell traces

Num sites	k-NN (%)	k-FP (%)	CUMUL (%)
10	91% ± 0.03	100% ± 0.00	99% ± 0.03
50	73% ± 0.01	91% ± 0.01	86% ± 0.03
100	68% ± 0.01	76% ± 0.02	76% ± 0.02
500	64% ± 0.00	72% ± 0.01	66% ± 0.01
1,000	59% ± 0.00	56% ± 0.01*	63% ± 0.01

Open World Onion Site Fingerprinting Results

- **One-class classification problem**
  - Site is the monitored site or other
  - We used a popular social networking site ( ) as the monitored site
  - Projection shows boundary that minimizes false positives
  - 80% of all errors were from 12 sites

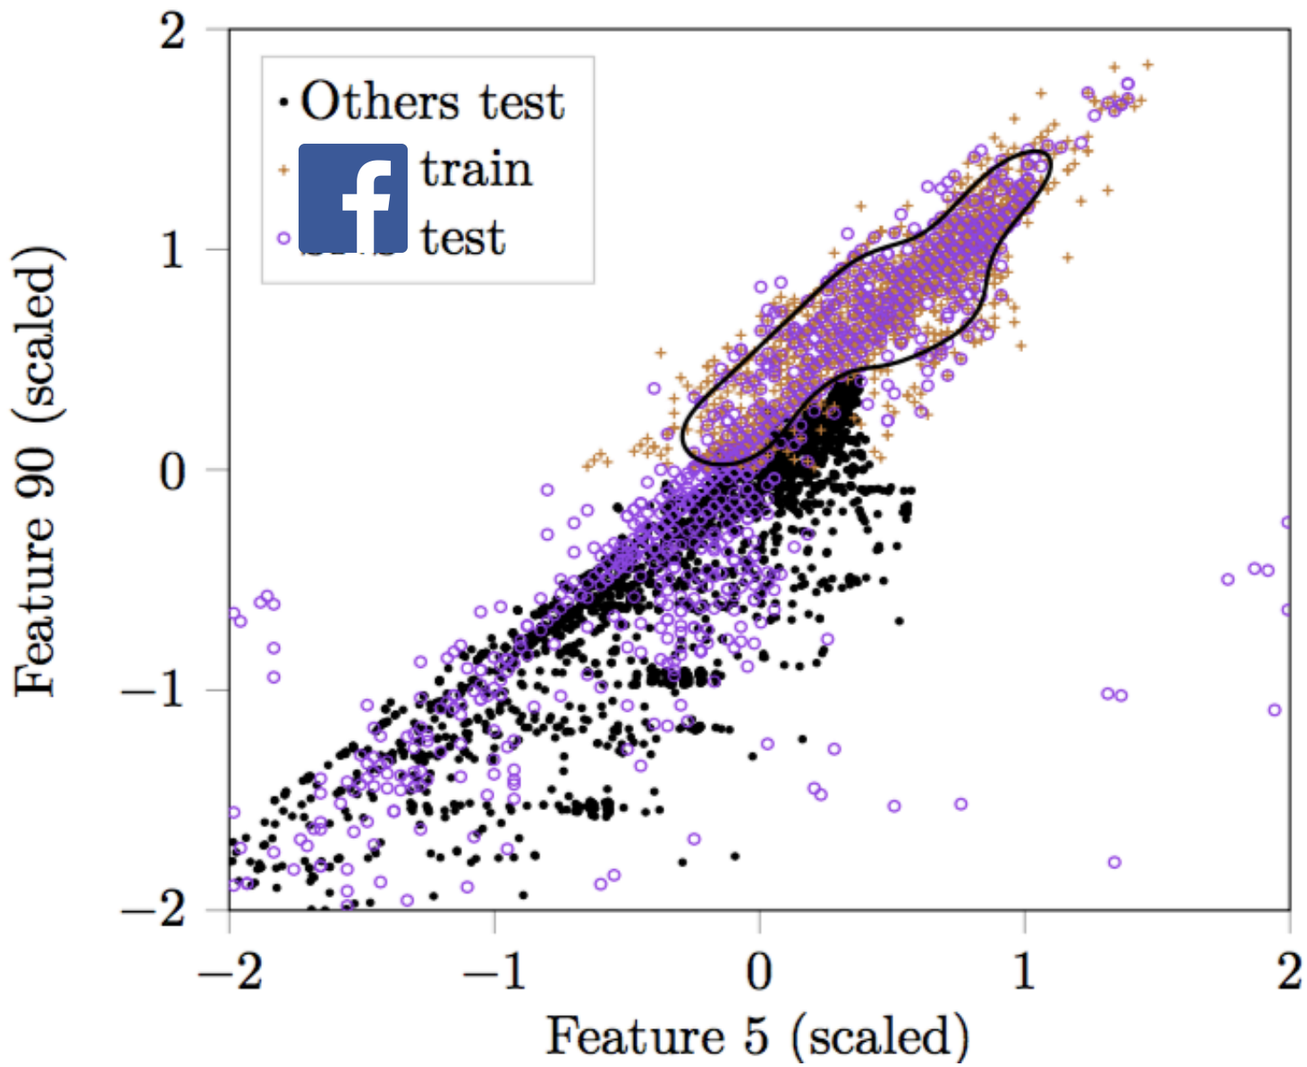

Open World Onion Site Fingerprinting Results

Base Rate Performance

- Precision is 50% at a base rate of 1%
- Precision decreases exponentially with the base rate

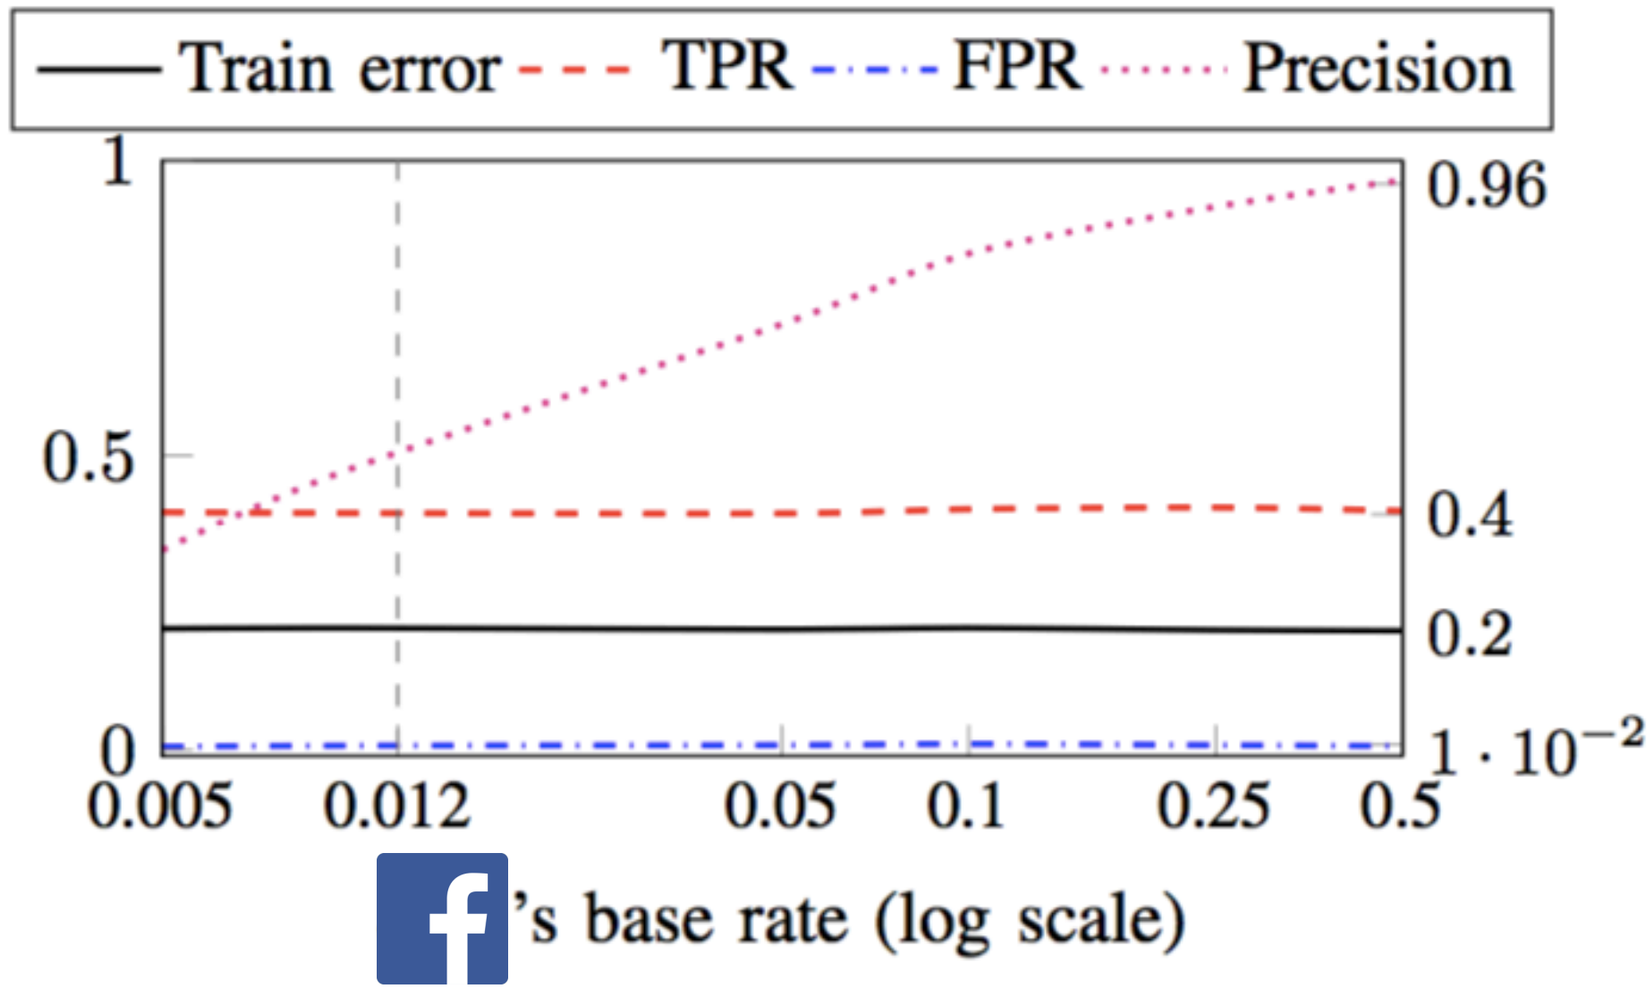

Onion Service Popularity Measurement

- Train classifiers on a social networking site front-page
- Apply trained classifiers to measure onion service popularity using privacy-preserving Tor measurement tool (PrivCount)

Classifying Circuits and Sites in Tor

- **Measured popular social network site that runs a single onion service**

- Enhanced PrivCount to classify circuit purpose, relay position, and site

- Three measurements:
  - Classify circuits from real Tor users
  - Classify circuits from ground truth crawler
  - Measure direct accesses to the ASN of (in the cases that we are the 3rd hop)

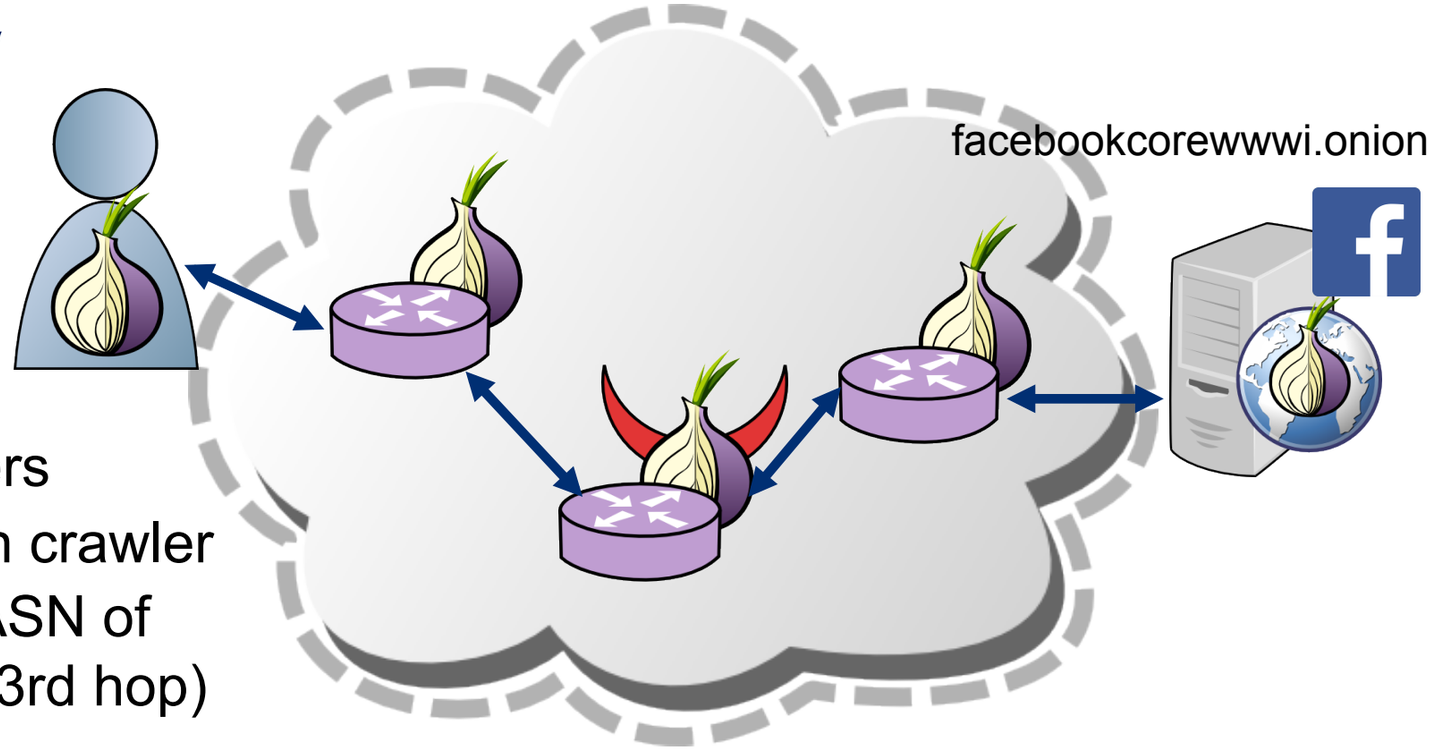

Classifying Circuits and Sites in Tor

- **Measured popular social network site that runs a single onion service**

- Enhanced P
  circuit purpos
  and site

- Three measu
  - Classify cir
  - Classify cir
  - Measure di
    (in the c

Ethical research:

- PrivCount provides differential privacy and secure aggregation of results
- No information is stored on disk
- Circuit-specific information is stored only for the life of the circuit (10 minutes)
- Consulted with Tor Research Safety Board to get feedback on methodology

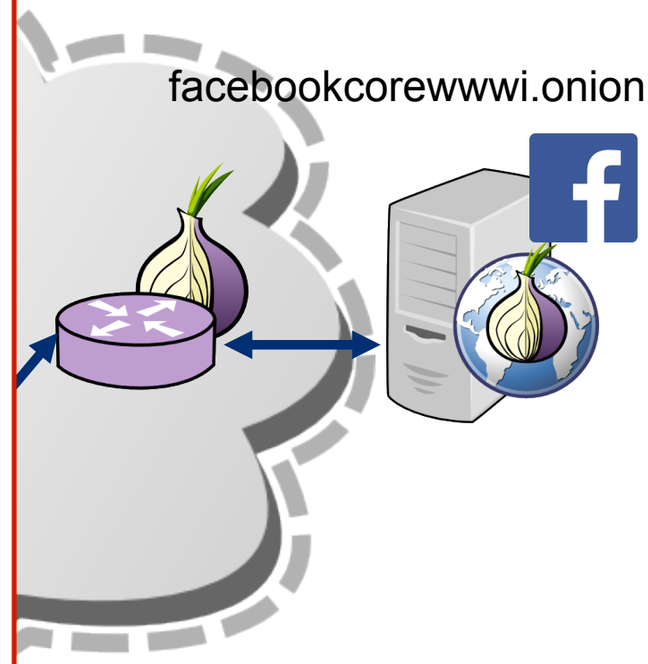

Classification Results

Crawler results (ground truth)

Classifier	True Positives	False Negatives
Purpose	100%	0%
Position	96.5%	3.4%
Site	60.0%	40.0%

Measurement pipeline results

Popularity	Direct	Classified
Purpose (onion service)	1.28%	4.48%
Site	0.52%	0.02%

Results include noise!

Conclusion

- Circuit and website fingerprinting is at least as accurate from middle relays as it is from the entry position

- The number of Facebook onion site visits was indistinguishable from noise

- More work needed to better understand middle relay threats

- All code is open-source:
  - github.com/onionpop
  - github.com/privcount
  - github.com/shadow

Contact:
Rob Jansen
U.S. Naval Research Laboratory
rob.g.jansen@nrl.navy.mil
robgjansen.com, @robgjansen

Onion Service Fingerprinting Classifiers

- **Train and test well known classifiers using packet and cell traces**

- k Nearest Neighbors (kNN) [Wang et al., 2014]
  - Averages over k closest instances according to Euclidean distance

- CUMUL [Panchenko et al., 2016]
  - Support vector machine (SVM) with radial basis function

- k-Fingerprinting (KFP) [Hayes and Danezis, 2016]
  - Random forest + kNN (with Hamming distance)