Website Fingerprinting at Internet Scale

Andriy Panchenko[1], Fabian Lanze[1], Andreas Zinnen[2],
Martin Henze[3], Jan Pennekamp[1], Klaus Wehrle[3], Thomas Engel[1]

[1]Interdisciplinary Centre for Security, Reliability and Trust (SnT), Luxembourg
[2]RheinMain University of Applied Sciences, Germany
[3]RWTH Aachen University, Germany

uni.lu
UNIVERSITÉ DU LUXEMBOURG

SnT
securityandtrust.lu

Background

- Why people use Tor...

- Privacy has become a general concern
- Access to the Internet is censored in many countries

Website Fingerprinting

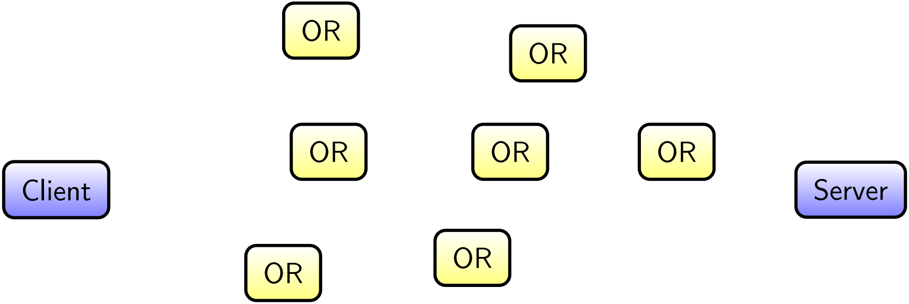

Tor: The Onion Router

- Most popular low-latency anonymization network
- Many users rely on Tor to access unfiltered information

Website Fingerprinting

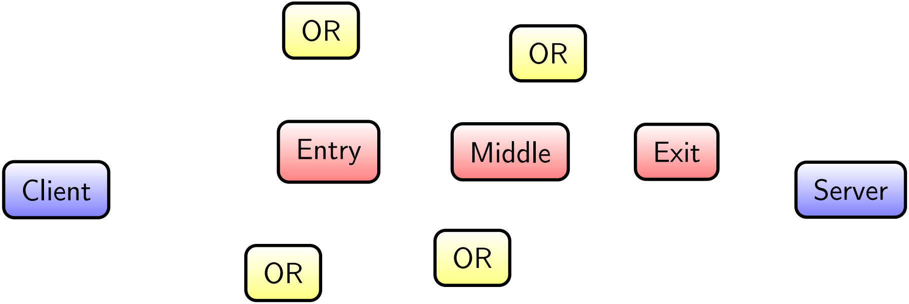

Tor: The Onion Router

- Most popular low-latency anonymization network
- Many users rely on Tor to access unfiltered information

Website Fingerprinting

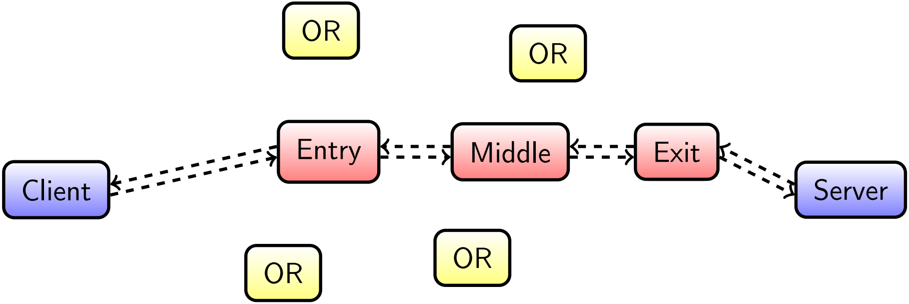

Tor: The Onion Router

- Most popular low-latency anonymization network
- Many users rely on Tor to access unfiltered information

Website Fingerprinting

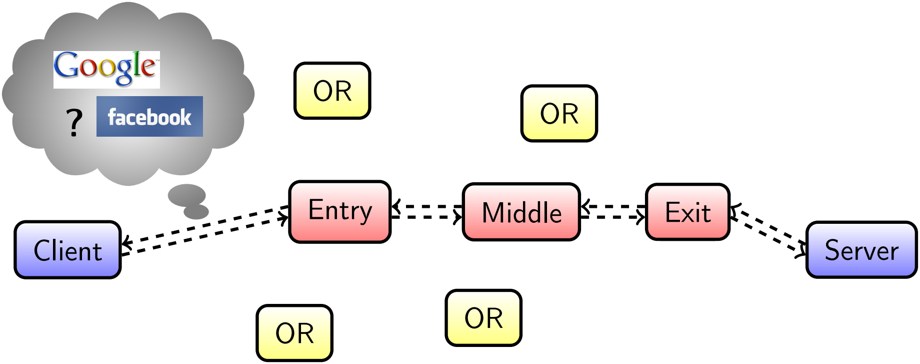

What is website fingerprinting?

- Identify website accessed without breaking cryptography
- Attacker is a *passive observer*
- Features based on packet size, direction, ordering, timing

Website Fingerprinting - state of the art

- Widely discussed and hot topic in anonymity research

State-of-the-art approach: **Wang et al.** (*Usenix Sec'14*)

- **k**-**N**earest **N**eighbor approach
- manually selected features (e.g., bursts, unique lengths)
- about 4,000 features
- recognition rates $> 90\%$

2 scenarios for evaluation

- **Closed world**: user visits only a fixed number of websites
- **Open world**: monitor set of sites (user may visit unknown sites)

Our method

Idea

- Don't try to guess which characteristics *may* be relevant
- Use a representation that *implicitly covers all* characteristics

Our feature set: $\Big(\underbrace{N_{\text{in}},N_{\text{out}},S_{\text{in}},S_{\text{out}},}_{\text{basic properties}}\quad\underbrace{C_1,\cdots,C_n}_{\text{cumulative features}}\quad\Big)$

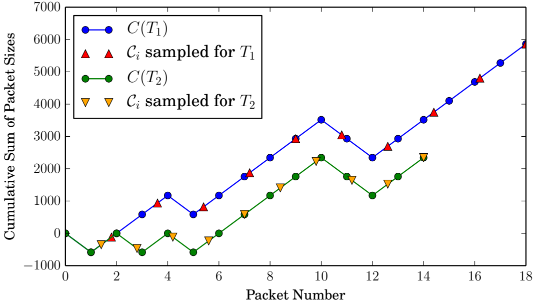

Example

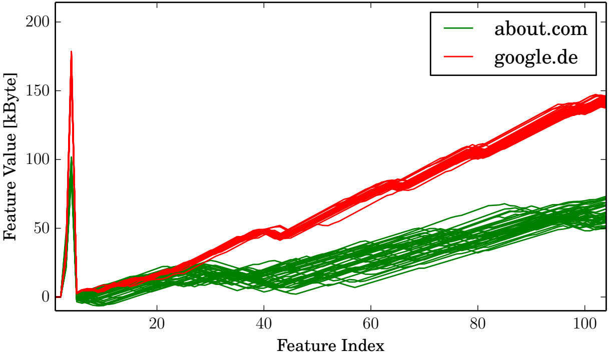

- **Fixed number** of distinctive characteristics from traces with **varying lengths**
- Fingerprints can be visualized
- Used as input for a Support Vector Machine

Layers of data representation

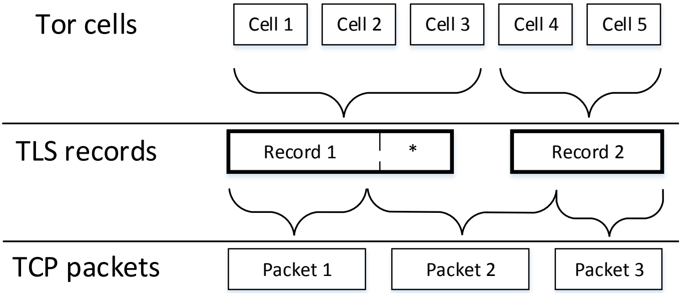

- Information src for feature extraction: Cell vs. TLS vs. TCP
- Practically nigligible effect on the classification accuracy

Comparison with state of the art – classification

Closed world

Accuracy [%] for *100 most popular websites*

	90 instances	40 instances
k-NN (3736 features)	90.84	89.19
Our method (104 features)	**91.38**	**92.03**

Open world

Foreground: 100 blocked websites, **background**: 9,000 popular websites

	TPR	FPR
k-NN	90.59	2.24
Our method	**96.92**	**1.98**

Comparison of computational performance

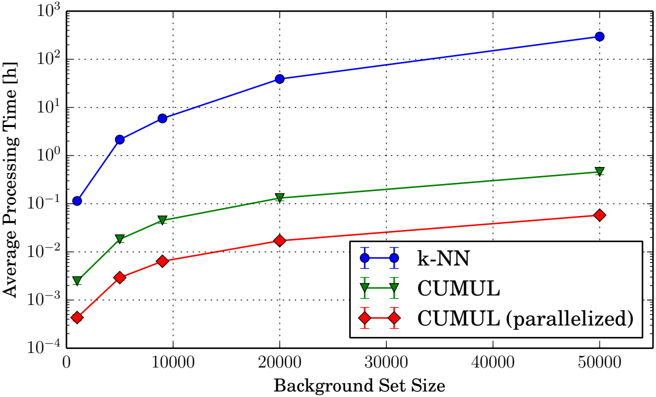

- Computation time for 100 random monitored pages in open world

Website fingerprinting in reality

Critique

- Data sets used are not representative!
    - too small, only popular websites / index pages
- Simplified assumptions, wrong metrics for evaluation

RND-WWW: How do people access the world wide web?

- Twitter
- Alexa-one-click
- Googling the trends
- Googling at random
- Censored in China

} > 120,000 web pages

Tor-Exit: Which pages do users actually access over Tor?

- Monitor a Tor Exit node ⇒ 211,148 web pages

Webpage fingerprinting at Internet scale

Question: *Does the attack scale under realistic assumptions?*

Which metric to evaluate?

- **Accuracy**: fraction of true results
- **True Positive rate / Recall**: fraction of monitored pages detected
- **False Positive Rate**: fraction of false alarms
    - **Problem**: misleading interpretation ⇒ *base rate fallacy*
- **Precision**: probability that the classifier is correct given it has detected a monitored page

Focus of evaluation

- Precision and recall for *increasing background set sizes*
- Random subset as foreground

Webpage fingerprinting at Internet scale

Question: *Does the attack scale under realistic assumptions?*

- Results for RND-WWW

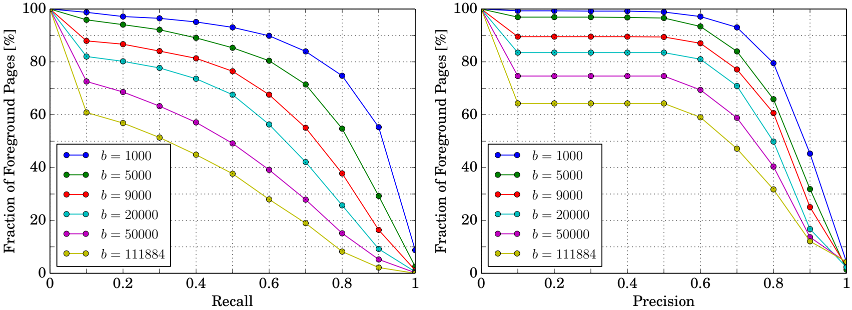

Webpage fingerprinting at Internet scale

Question: *Does the attack scale under realistic assumptions?*

- Results for Tor-Exit

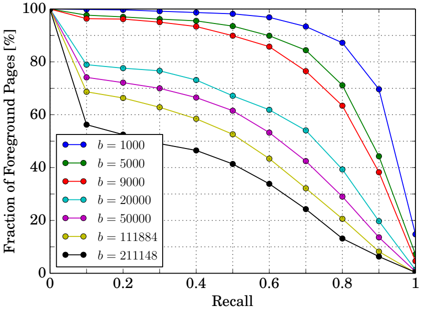

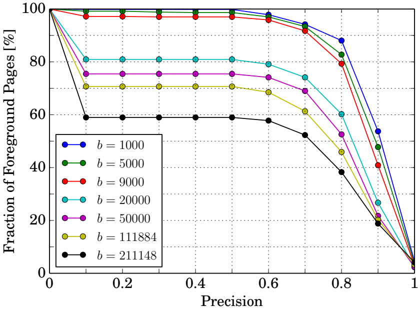

Webpage fingerprinting at Internet scale

Question: *Does the attack scale under realistic assumptions?*

- Results for Tor-Exit

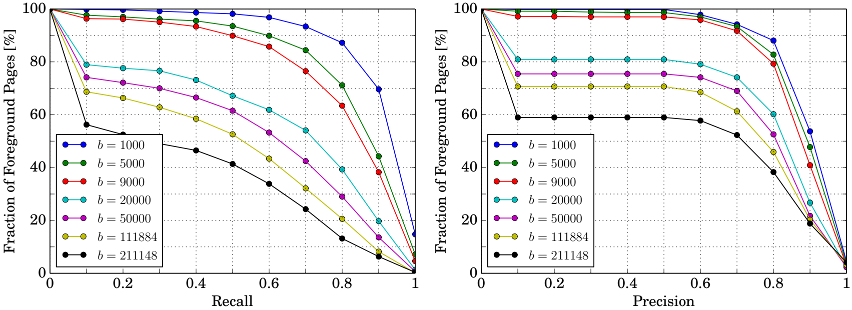

Answer: No.

Webpage fingerprinting at Internet scale

Question: *Is it at least possible for* ***certain*** *pages?*

Webpage fingerprinting at Internet scale

Question: *Is it at least possible for **certain** pages?*

- Minimum number of mistakenly confused pages

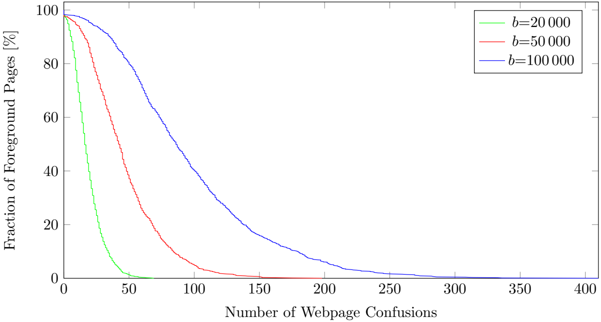

No single page without a confusingly similar page in a realistic universe.

How about fingerprinting web***sites***? (1/2)

- A website is a **collection of web pages** served under the same domain
- Is it possible to fingerprint a website when *only a subset of its pages* are available for training?

Experiment: 20 websites

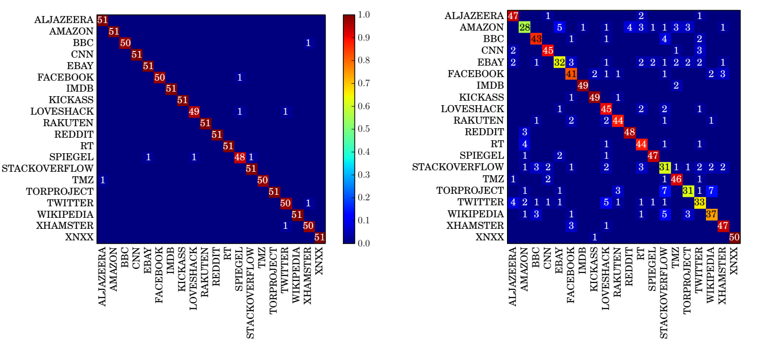

(a) only index pages

(b) different pages

How about fingerprinting web***sites***? (2/2)

- Transition of results from closed-world to the realistic open-world setting is typically not trivial
- Website fingerprinting **scales better** than webpage fingerprinting

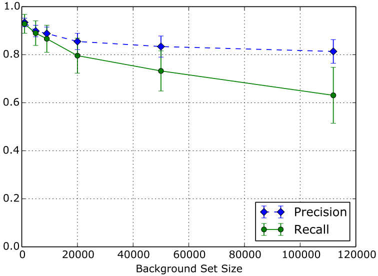

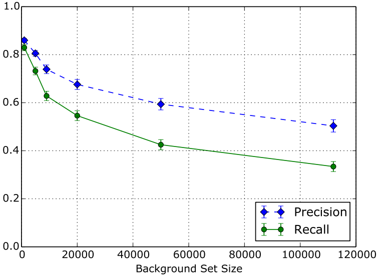

Summary

- Our classifier with 104 features **outperforms** state of the art
- Alarming results under simplified assumptions **can't be generalized**
- Webpage fingerprinting **does not scale** for appropriate universe sizes for any webpage
- Website fingerprinting is not only more realistic and also significantly more effective
- *Conclusions drawn need to be reconsidered*

Scripts and RND-WWW dataset:
`http://lorre.uni.lu/~andriy/zwiebelfreunde/`

We are hiring!

uni.lu

UNIVERSITÉ DU
LUXEMBOURG

Our lab within the Interdisciplinary Centre for Security, Reliability and Trust (Uni Luxembourg) is looking for PhD candidates and PostDocs in the area of anonymity and privacy

More information: `http://secan-lab.uni.lu/jobs`