DNS-based User Tracking (Attacks and Defenses)

Presenter: Zhou Li

Data-driven Security and Privacy (DSP) Lab
EECS department, University of California, Irvine
DINR'23, 02/22/2023

Outline

Background

Threat Model

Attack: DSCorr

Defense: LDPResolve

Conclusion

Paper published: “Hide and Seek: Revisiting DNS-based User Tracking. Deliang Chang, Joann Qiongna Chen, Zhou Li, and Xing Li. EuroSP’22”.

Background: DNS-based User Tracking

- Users send DNS queries before almost every network activities.
- Different users have different preferences.
- Can we track a user by their DNS queries?
  - Privacy violation

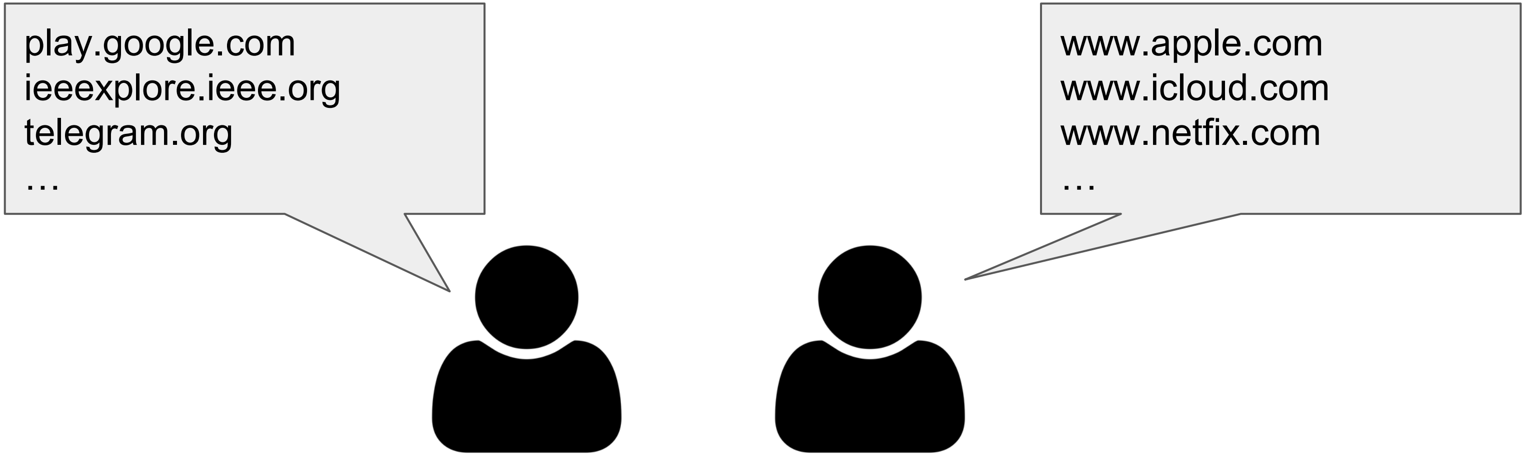

Attack: Threat Model

- Goal: track users based on their DNS queries
  - E.g., public recursive resolvers
- Challenge: a user’s identifier, aka, source IP keeps changing
  - E.g., DHCP, moving from one access point to another, cellular network
- This is an inference/classification problem
  - Attacker’s input: **Session,** DNS queries from one source IP in a time windos
  - Attacker’s output: **user ID** (real or pseudo)

Threat Model

- Formalization of DNS-based user tracking
  - Link different sessions of a same user from different source IPs.

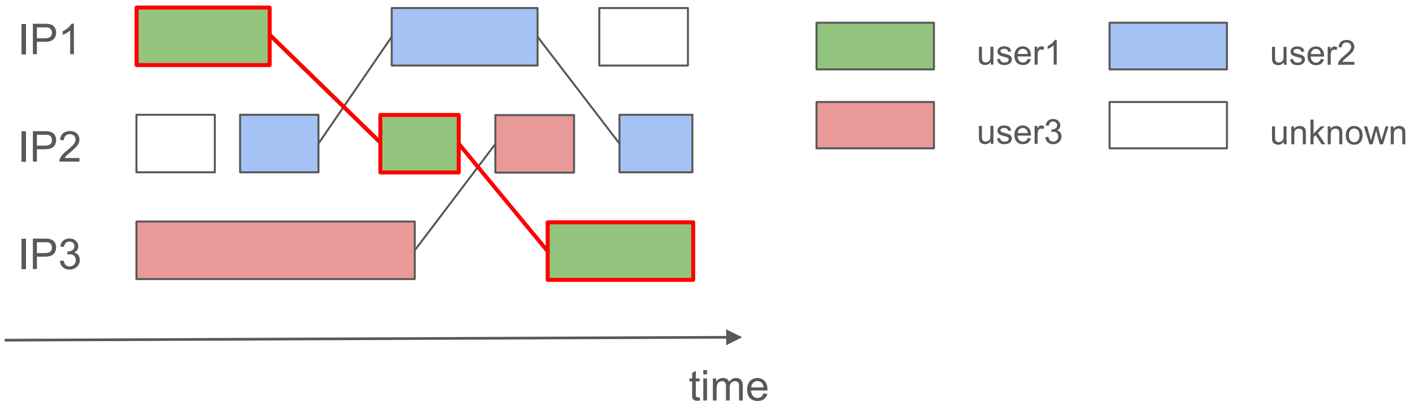

Existing Attacks

- Supervised, semi-supervised or unsupervised learning
  - Feature extraction from DNS queries
  - Bayesian classifier, KNN, Dirichlet multinomial mixture
  - Fixed threshold
- All assuming a closed-world setting
  - The attacker already knows the set of users before tracking
- How about open-world setting?
  - Unknown user can be encountered during tracking

[1] Dominik Herrmann, Christian Banse, and Hannes Federrath. Behavior-based tracking: Exploiting characteristic patterns in dns traffic. Computers & Security, 39:17–33, 2013.
[2] Dominik Herrmann, Matthias Kirchler, Jens Lindemann, and Marius Kloft. Behavior-based tracking of internet users with semi-supervised learning. In 2016 14th Annual Conference on Privacy, Security and Trust (PST), pages 596–599. IEEE, 2016.
[3] Dae Wook Kim and Junjie Zhang. You are how you query: Deriving behavioral fingerprints from dns traffic. In International Conference on Security and Privacy in Communication Systems, pages 348–366. Springer, 2015.
...

Our Attack: DSCorr

STEP 1	Convert domains to domain embedding vectors.
STEP 2	Build user profiles (clusters of sessions) from a labeled session set.
STEP 3	Given an unlabeled session ***s***, identify ***k*** nearest labeled clusters through a data-sketching process.
STEP 4	Compute the nearest distance between session ***s*** and ***k*** clusters under user-centirc threshold for open-world setting.

Domain Embedding

- Domain distance: 0 or 1 by previous works
  - Too coarse-grained
- Fine-grained domain distance based on domain context
  - Domains usually visited together should have small distance
- Use Word2Vec (NLP) to build domain embedding vectors
  - Domain -> Word
  - DNS session -> Context

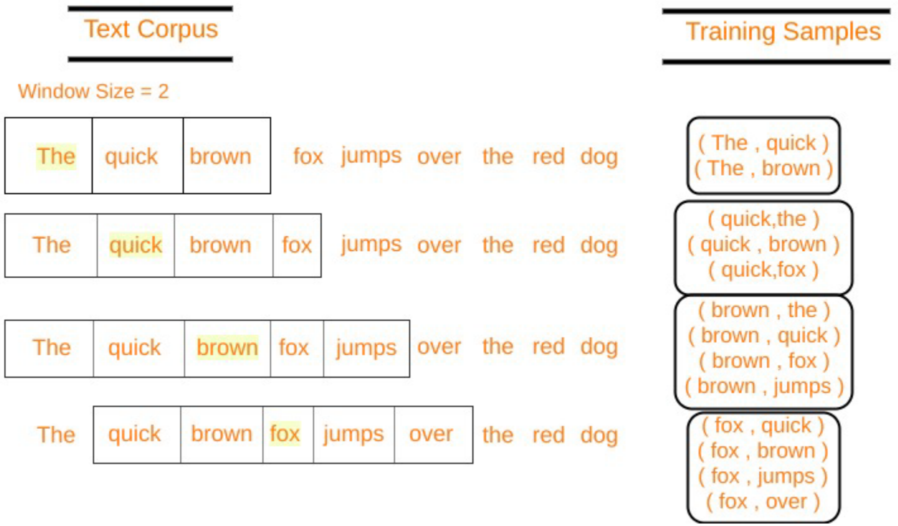

SkipGram of Word2Vec

Evaluation of DSCorr

- Different tracking methods: Jaccard/Cosine/Bayesian Classifier/DSCorr
- Different feature: unigram & bi-gram
- Different number of sessions in labeled set for each user

#	jac	cos	bay	ja-bi	co-bi	ba-bi	DSCorr
1	42.2	40.7	37.4	45.4	40.1	36.5	**52.6**
2	56.0	52.8	54.8	59.2	52.8	54.3	**67.5**
3	67.2	60.3	65.7	67.2	60.3	65.8	**74.4**
5	74.8	69.3	76.3	74.8	69.3	76.8	**80.5**
10	78.8	78.0	86.2	82.7	77.6	87.3	**87.4**

Table. Tracking accuracy under closed-world setting.

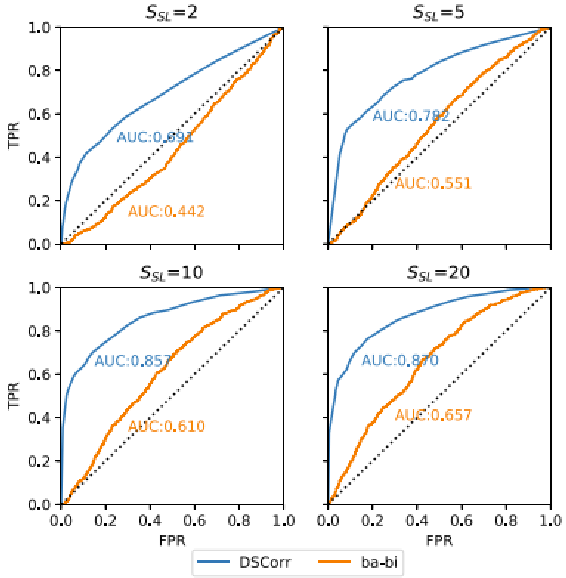

Fig. Tracking accuracy under **open-world setting**.

Findings:

- DSCorr is more effective under closed-world setting, especially when there's **less labeled data**.
- Auto-threshold works. It allows DSCorr to work under open-world setting.
- Popular domains affect user tracking.

Defense: Local Differential Privacy (LDP)

- The data collector is untrustworthy
- Noises added to the clients' data before collection
- LDP guarantees the information leakage after noises are bounded by ϵ
- Used by Apple to collect emoji usage …

Definition 1 (ϵ-Local Differential Privacy [89]). *An algorithm $\mathcal{A}$ satisfies ϵ-local differential privacy (ϵ-LDP), where $\epsilon > 0$, if and only if for any pair of input x_1 and x_2, we have*

$$\forall y \in \mathrm{Range}(\mathcal{A}) : \frac{\Pr[\mathcal{A}(x_1) = y]}{\Pr\left[\mathcal{A}\left(x_2\right) = y\right]} \leq e^{\epsilon} \qquad (1)$$

where $\mathrm{Range}(\mathcal{A})$ *denotes the set of all possible output results of an algorithm* $\mathcal{A}$.

Our Defense Method: LDPResolve

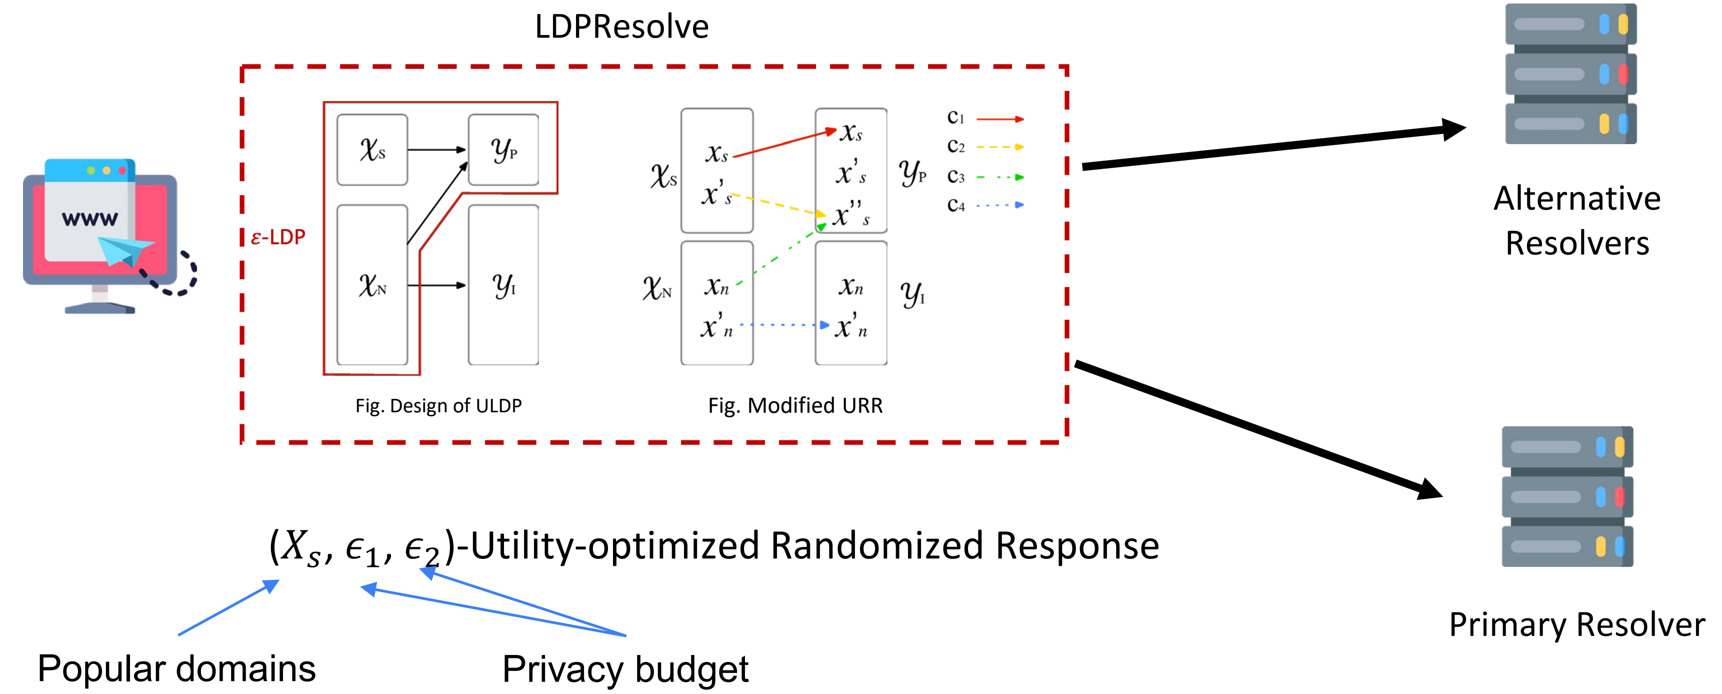

Murakami, Takao, and Yusuke Kawamoto. Utility-Optimized Local Differential Privacy Mechanisms for Distribution Estimation. USENIX 2019.

Design of LDPResolve

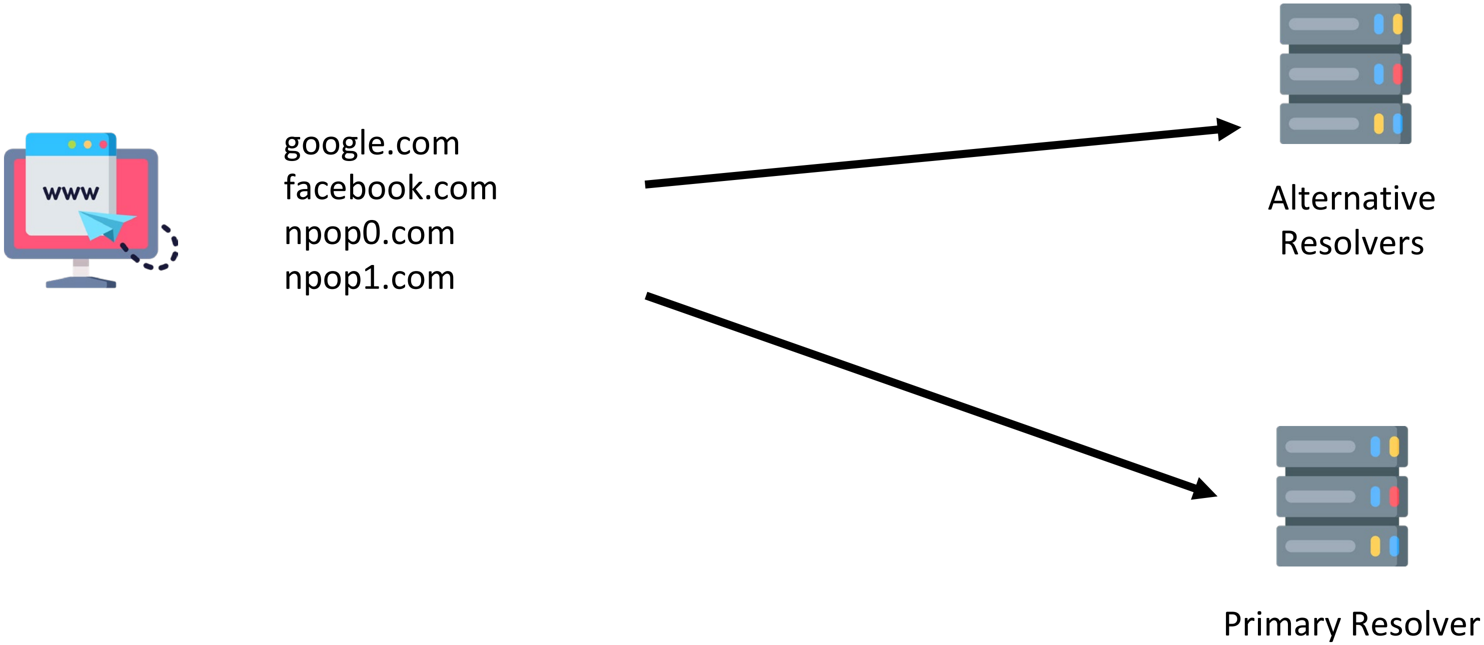

Design of LDPResolve

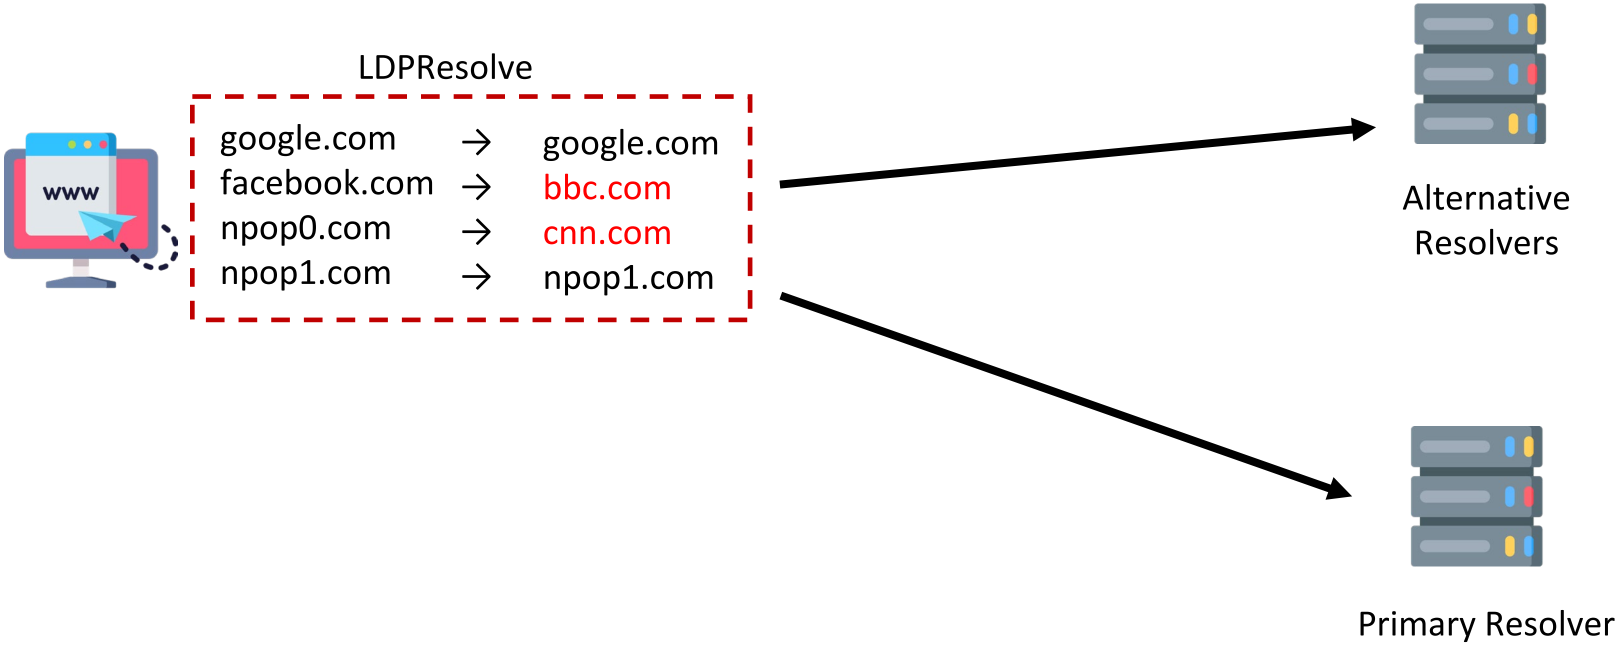

Murakami, Takao, and Yusuke Kawamoto. Utility-Optimized Local Differential Privacy Mechanisms for Distribution Estimation. USENIX 2019.

Design of LDPResolve

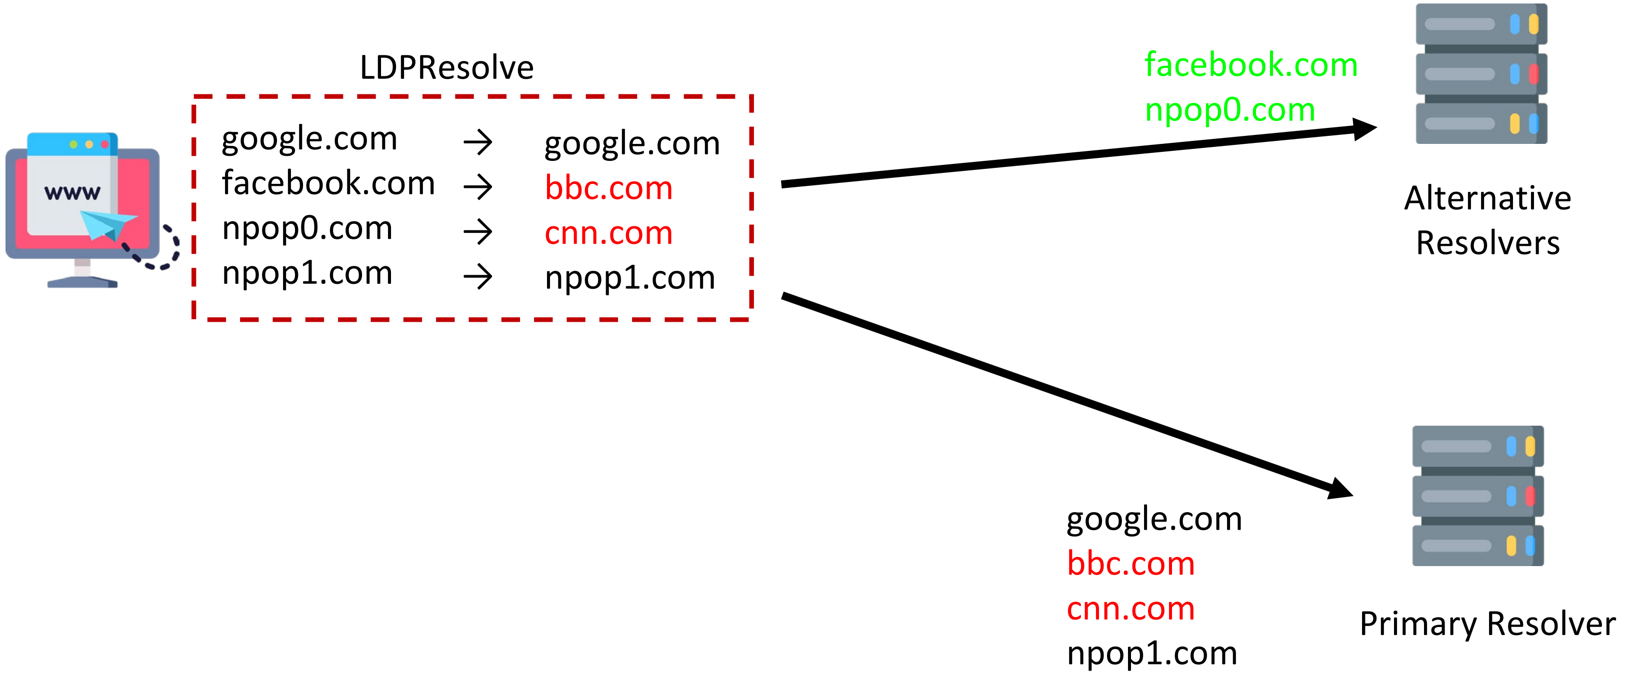

Murakami, Takao, and Yusuke Kawamoto. Utility-Optimized Local Differential Privacy Mechanisms for Distribution Estimation. USENIX 2019.

Evaluation of LDPResolve: Privacy

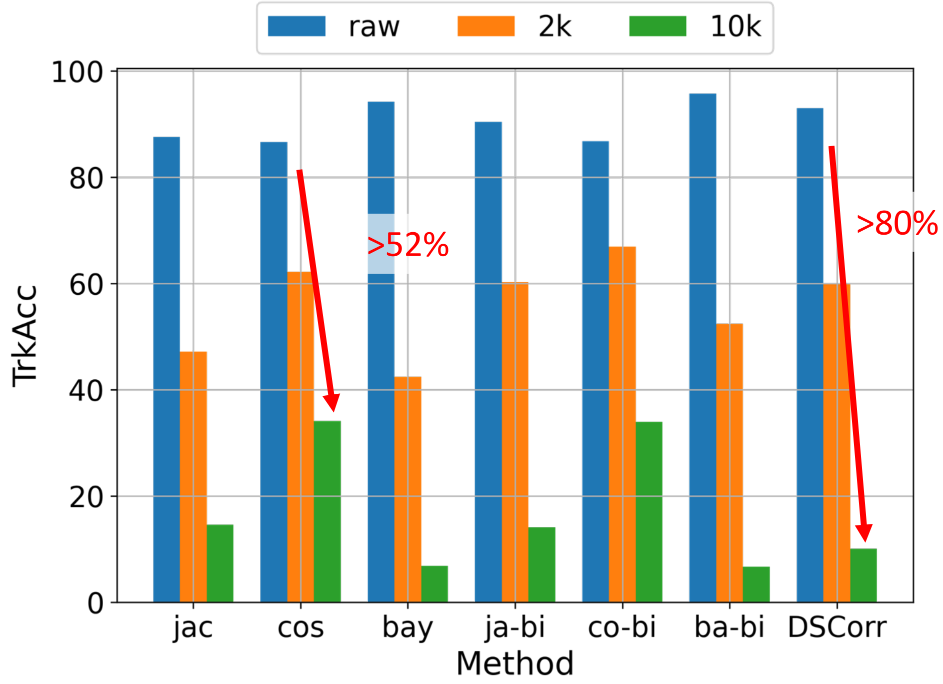

Fig. Tracking Accuracy given different sensitive set size (i.e., 2k and 10k)

Evaluation of LDPResolve: Utility

ε_1	TrkAcc	std	std_s	std_n
15	38.7	332.30	1279.53	3.48
10	34.1	343.66	1279.63	5.71
9	28.4	352.52	1279.94	6.85
8	19.5	360.62	1280.39	8.54
7	10.2	365.38	1279.76	10.66
6	3.7	367.61	1279.92	10.75
5	1.4	368.45	1279.01	11.20
2	0.2	369.12	1280.31	10.84

ε_2	TrkAcc	std	std_s	std_n
10	84.8	121.59	241.10	3.27
8	80.2	264.24	731.94	3.80
7	70.3	305.47	967.65	5.31
6	57.4	326.82	1127.27	5.52
5	43.6	336.81	1214.65	5.67
2	34.1	343.95	1279.63	5.71
0.5	33.9	343.95	1282.55	4.38

N_s	TrkAcc	std	std_s	std_n
1k	68.0	363.13	2552.22	1.72
2k	62.2	388.23	2205.18	2.15
5k	48.8	376.73	1669.81	6.54
10k	34.1	343.66	1279.63	5.71
20k	23.3	304.17	949.84	7.13

Key Terms

N_s: Size of sensitive set

ε_1: Overall privacy Budget

ε_2: Privacy Budget for sensitive domains. $\varepsilon_2 \leq \varepsilon_1$

Conclusion

- DNS-based user tracking is a real privacy concern
  - Existing works are effective under closed-world setting.
  - Our attack DSCorr is effective in both closed-world and open-world settings..
- Popular domain is the key to DNS-based user tracking.
- LDPResolve could be effective in terms of defeating tracking.
  - LDP ensures the privacy leakage is bounded regardless of the attack methods

Sponsors

- AMP

- EAGER SaTC
- IMR
- **CAREER**

This talk

Team (DSP Lab)

Zhou Li

Joann Chen

Jiacen Xu

Qifan Zhang

Xuesong Bai

Xiang Li