AdVeil: Private Targeted Advertising

Sacha Servan-Schreiber & Kyle Hogan with Srini Devadas



Overview of this talk

- 1. Background on targeted online advertising.
- 2. Guarantees, threat model, and assumptions.
- 3. System overview and performance.
- 4. Details on private targeting and fraud prevention.
- 5. Intersection attacks.
- 6. Browser integration.

Not covered (see paper):

1. Nitty-gritty system and implementation details.

Background







Advertising Pipeline



Ads originate at the "Advertiser"

In this case, Patagonia



The "Broker" manages ad campaigns and user data In this case, Google



The "Publishers" are websites that display ads to users and send targeting data back to the Broker



Follow the \$\$\$:

● Advertisers ⇒ Broker for targeting and delivery



Follow the \$\$\$:

- Advertisers ⇒ Broker for targeting and delivery
- **Broker** ⇒ **Publishers** for displaying ads



Follow the \$\$\$:

- Advertisers ⇒ Broker for targeting and delivery
- **Broker** ⇒ **Publishers** for displaying ads
- ⇒ Advertisers make more money

System Overview

Generic advertising ecosystem



AdVeil guarantees:

Broker wants: correctness (fraud prevention). **Users want:** privacy.

• Correctness

- 1. Clients see ads related to their interests.
- 2. All ads that **honest** clients interact with are billed correctly.
- 3. All ads that **malicious (bot)** clients interact with are discarded.

• Privacy

- 1. **Unlinkability:** between ads reported on and set of users.
- Security
 - 1. Fraud prevention cannot be exploited to identify clients.*
 - 2. Fraud prevention is covert.

AdVeil threat model:

Adversary model: Broker is malicious-but-rational.

- The **Broker** and **advertisers** want to link users to ads.
- The **clients** and **publishers** collude to perpetrate reporting fraud.

By *rational* we specifically mean that:

• The **Broker** is monetarily incentivized to maintain correct functionality of AdVeil.

Designing around rational adversary behavior.

Broker's requirements:

- 1. **Billing Integrity**: fraud prevention mechanism is resilient to large-scale botnets.
- 2. Covert fraud detection (avoid alerting bots [KLOR'20])

Using fraud prevention mechanism to link a user to an ad either:

1. Compromises billing integrity *(irrational)*

OR

2. Loses covertness of bot detection *(irrational)*.

Designing around rational adversary behavior.

Broker's requirements:

- 1. **Billing Integrity**: fraud prevention mechanism is resilient to large-scale botnets.
- 2. Covert fraud detection (avoid alerting bots [KLOR'20])

Using fraud prevention mechanism to link a user to an ad either:

1. Compromises billing integrity (irrational)

OR

2. Loses covertness of bot detection *(irrational)*.

A rational Broker will use fraud detection correctly.





Targeting







Targeting

Targeting

1. User profile constructed locally by the client.

Browsing history, shopping history, and other local data.

2. Client gets ad IDs for targeted ads (protocol described later).

- Ad IDs are small targeting identifiers later used to retrieve ads.
- Executed *periodically* (e.g., once a day) by the client.

3. Also gets signed tokens for verified reporting:

Broker issues a token based on client's IP address (human/bot).



Delivery

Delivery

- 1. Client retrieves ads corresponding targeted ad IDs.
 - Performed *on-demand* through an anonymizing proxy
- 2. Broker sends a new token for verified reporting.
 - Token binds the client to only report on the retrieved ad.
- 3. Retrieved ad(s) get displayed by the user's client on publishers.



Reporting

Reporting

1. User interactions generate reports, for example:

(ad ID: impression, ad ID: click).

2. Client sends report and tokens on-demand.

- Performed through an anonymizing proxy.
- Targeting token tells the Broker if client is a human/bot.
- Delivery token tells the Broker if client reported on the correct ad.
- 3. Broker validates each token and discards invalid reports.



Targeting





(e.g., Tor, VPN0 or I2P)

Timing of AdVeil





Performance

(TL;DR AdVeil is pretty fast in practice)



Takeaways

@ 1M ads and 16 LSH hash tables* we get:

- Good targeting accuracy (> 95%) in practice.
- Reasonable latency (a few seconds) even on low-bandwidth clients (e.g., mobile devices).

Technical notes

- Targeting only needs to be performed periodically (e.g., daily) and offline (without user involvement).
- Bottleneck is processing time of PIR (SealPIR).
- Server throughput (clients / second) is much faster compared to latency (see paper for full evaluation).
- Batching ads amortizes server *processing* time when targeting k ads simultaneously.



Takeaways

- Ads are much larger than ad IDs, making PIR (even two server*) far too slow for delivery!
- Tor can support delivery of video ads quickly, while also hiding client identity.

Technical notes

- Delivery can support on-demand ad selection logic where the Broker picks an ad in real-time for the provided ad ID.
- Any anonymizing proxy providing sender anonymity is enough (e.g., VPNO and I2P).

Technical Details

Technical Details

- Targeting
- Fraud-prevention
- See paper for Delivery and Reporting.

Targeting details

Main idea

- 1. Build a targeting data structure using LSH.
- 2. Client queries for neighbors (ad IDs) using PIR.

LSH+PIR = private targeting mechanism.

- 3. For fraud prevention:
 - Use anonymous one-time-use token.



Targeting (with no privacy)

Abstractly: Nearest Neighbor Search



Targeting (with no privacy)

Abstractly: Nearest Neighbor Search

Hash Table



Targeting (with no privacy)





Private Targeting



Profile 1000100001

















• Use one-time-use anonymous tokens with private metadata [KLOR'20]

Fraud Prevention

 \circ Generalization of Privacy Pass [DGS'18] used by CloudFlare in the wild



• Use one-time-use anonymous tokens with private metadata [KLOR'20]

Fraud Prevention

• Generalization of Privacy Pass [DGS'18] used by CloudFlare in the wild



Fraud prevention details

- Use one-time-use anonymous tokens with private metadata [KLOR'20]
 - Generalization of Privacy Pass [DGS'18] used by CloudFlare in the wild
- Encode one bit b of private metadata: human(b=1) OR bot(b=0)
 - Client cannot read the private metadata bit b



A technical problem

Observation: Implicit third type of token!

Broker can use *invalid* tokens to tag users.



A technical problem

Observation: Implicit third type of token!

Broker can use *invalid* tokens to tag users.



Solution: only have "valid" and "invalid" tokens.

Then at most 2 types of tokens and corresponding anonymity sets.



Token Abuse: Irrational Behavior

Broker either (1) loses covert bot detection by refusing to serve bots* OR(2) loses billing report accuracy by mixing the reports of bots and honest clients

Both (1) and (2) result in a direct financial implications \Rightarrow irrational behavior.



Leakage: Intersection Attacks

For every advertising period the Broker learns which **users** participated and what **ads** were shown



However, it learns **nothing** about the mapping between users and ads

Leakage: Intersection Attacks

The set of users may change across advertising periods, and users are inherently correlated with the ads they see.



Period n

Period *n*+1

Leakage: Intersection Attacks

The set of users may change across advertising periods, and users are inherently correlated with the ads they see.



Solving intersection attacks

Not possible to completely eliminate intersection attacks without requiring:

- all users to participate in every period

OR

- using a fixed set of ads per period

Leakage is minimized when:

- 1. Ads are less personalized (i.e., a large number of people could see a given ad)
- 2. Churn is ~50% (i.e., a large group of people is staying/joining/leaving)

Integration with browsers

We use Tor

- Any other anonymous proxy is also a possibility
 - VPN0 [VANPPL'19]
 - I2P [ZH'11]

Lightweight on the client

- Targeting protocol requires ~12MB of communication; **comparable to visiting 4 websites**.
- Processing time extremely light: 300 microseconds (on a single core).

Potential roadblocks

- Requires anonymous proxy to be "built into" the browser (as done in Brave already).
- Requires many more servers to process targeting,

e.g., between 100 and 1,000 servers with **10M users** being targeted weekly.

Thank you!

Sacha Servan-Schreiber (3s@mit.edu)

Kyle Hogan (<u>klhogan@mit.edu</u>)

Paper + prototype code available at: <u>https://adveil.com</u>.