AdVeil: A Private Targeted-Advertising Ecosystem

Sacha Servan-Schreiber
MIT CSAIL
3s@mit.edu

Kyle Hogan
MIT CSAIL
klhogan@mit.edu

Srinivas Devadas
MIT CSAIL
devadas@mit.edu

ABSTRACT

This paper presents AdVeil, a privacy-preserving advertising ecosystem with formal guarantees for end users.

AdVeil is built around an untrusted advertising network which is responsible for brokering the display of advertisement to users. This ad network targets relevant ads to users *without* learning any of the users' personal information in the process. Our targeting protocol combines private information retrieval with standard, locality-sensitive hashing based techniques for nearest neighbor search. By running ad targeting in this way, users of AdVeil have full control over and transparency into the contents of their targeting profile.

AdVeil additionally supports private metrics for ad interactions, allowing the ad network to correctly charge advertisers and pay websites for publishing ads. This is done *without* the ad network learning which user interacted with an ad, only that some honest user did. AdVeil achieves this using an anonymizing proxy (e.g., Tor) to transit batched user reports along with unlinkable anonymous tokens with metadata to certify the authenticity of each report.

We build a prototype implementation of AdVeil which we evaluate on a range of parameters to demonstrate the applicability of AdVeil to a real-world deployment. Our evaluation shows that AdVeil scales to ad networks with millions of ads, using state-of-the-art single-server private information retrieval. A selection of ads from a database of 1 million ads can be targeted to a user in approximately 10 seconds with a single 32-core server, and can be parallelized further with more servers. Targeting is performed out-of-band (e.g., on a weekly basis) while ad delivery happens in real time as users browse the web. Verifying report validity (for fraud prevention) requires less than 300 microseconds of server computation per report.

1 INTRODUCTION

Internet advertising is a $124 billion industry that relies on pervasive tracking of internet users for the purpose of serving them relevant advertisements. This process is known as *targeted* advertising and has been the focus of recent controversy due to the often highly personal nature of the data being used and the invasiveness of the collection practices [41, 72, 80]. Proposals for moving away from targeted advertising often tout *contextual* advertising as a privacy focused alternative [51, 61, 78]. In contextual advertising, ads are chosen based *only* on the website they will be displayed on, and *not* on personal data, eliminating the need for web tracking.

However, relevancy of advertisements for end users is believed to increase user engagement which, consequently, increases profit [55, 67]. Google and Facebook, the most prominent ad-tech companies today, respectively earned 83% and 99% of their 2019 revenue from advertising alone [37, 48, 88]. Because of this, ad networks have proven unwilling to move away from targeted advertising and the associated web tracking. Recent proposals from Google for a privacy preserving alternative [96, 97] fall short, with privacy advocates such as the Electronic Frontier Foundation (EFF) referring to them as being "the opposite of privacy-preserving technology" [29].

Other solutions for privacy-preserving targeted advertising have practical performance limitations or fail to achieve targeting accuracy comparable to non-private advertising [9, 14, 43, 45–47, 52, 73, 90, 91, 96, 97]. AdVeil addresses these problems.

AdVeil is a low-latency, scalable solution for targeted advertising with clear-cut privacy guarantees. The main goal of AdVeil is to provide *unlinkability* between users and their personal data. We define *personal data* to mean any information about the user, including which ads they interact with. The interests and demographics contained in users' targeting profiles are not sent in the clear – even anonymized – to the ad network as part of the targeting process. However, due to the nature of targeting, knowledge of the ads a user sees can be used to infer the, clearly personal, data that makes up the user's profile. As such, AdVeil ensures that the ad network learns **only** which ads are displayed to (and clicked by) users, but not which user saw any given ad.

AdVeil is fully compliant with data privacy legislation such as General Data Protection Regulation (GDPR) [2] and California Consumer Privacy Act (CCPA) [3]. Users' profiles are held locally, with the users themselves having full control over and transparency into *which* of their features are used for targeting. They may even opt out of targeted advertising entirely, in which case websites can display contextual ads that are related only to the page itself, not to the users viewing them.

Giving users control over their own data is important for reasons that stem beyond privacy. Companies have been found to use protected demographics like gender, race, or religion to target advertisements in a discriminatory manner [5, 22, 49, 98]. This practice can limit the visibility of needed services, such as home loans or continuing education, to people who might have benefited from them [33, 74].

Finally, addressing ad fraud is a crucial requirement for online advertising [83, 105]. Malicious parties may attempt to generate false ad interactions for the purposes of artificially running down an advertiser's budget or increasing revenue for a website for displaying ads. AdVeil provides an integrated fraud-prevention mechanism for detecting bots and ensuring accuracy of reports.

AdVeil is realized using well established building blocks which we combine to provide a scalable, accurate, and privacy-preserving advertising ecosystem. We use private information retrieval in conjunction with standard nearest neighbor search data structures to achieve private ad targeting. Our fraud-prevention mechanism is built using unlinkable tokens with metadata [56, 79]. Finally, we use any common anonymizing proxy (e.g., Tor [35]) to hide the identities of users from the ad network during ad delivery and reporting.

To the best of our knowledge, we are the first to apply private-information-retrieval to privately *target* ads to users. We believe that this technique may be of independent interest for related problems such as privacy-preserving recommendation systems.

Contributions. AdVeil makes the following contributions:

(1) A novel and efficient protocol for ad targeting that reveals no information about the user profile. Our protocol may have applications to other recommendation systems that require user privacy.
(2) An integrated fraud prevention mechanism based on anonymous tokens. Our solution remains fully compatible with private browsing and existing ad-fraud counter-measures.
(3) Out-of-the-box compliance with recent privacy regulations such as GDPR and CCPA, in addition to full compatibility with existing anti-tracking and privacy-enhancing technologies.
(4) A prototype implementation which we empirically evaluate to demonstrate applicability to a real-world deployment. Our code is open-source and available at http://adveil.com/code.

Limitations. We highlight several drawbacks of AdVeil compared to the status-quo in (non-private) targeted online advertising.

(1) AdVeil imposes a computational overhead on the ad network, which translates to higher operational costs. However, in our evaluation (see Section 7) we provide concrete cost estimates and show that AdVeil remains profitable, even if deployed on off-the-shelf cloud infrastructure.
(2) AdVeil is intended to *disincentivize* web tracking by ad networks and is fully compatible with the use of any existing web tracking defenses. However, it is not itself a defense against browser fingerprinting or other tracking mechanisms.
(3) AdVeil requires cooperation of browsers. We note, however, that major browser vendors are already providing local support for several aspects of privacy-preserving advertising [73, 96, 99].

2 BACKGROUND

In this section we start by introducing the existing, non-private advertising ecosystem. This ecosystem consists of a set of parties who engage in different stages of an *advertising pipeline*.

This pipeline supports many different styles of targeted advertising, each with tradeoffs between their data requirements and the relevancy of ads delivered to users. We provide details on a selection of prevalent advertising styles in Section 2.3.

2.1 Participants

We use standard terminology for parties that comprise the advertising pipeline [9, 43, 45, 52, 90]:

- **Users**: people browsing the web and viewing ads;
- **Clients**: web browsers controlled by users (e.g., Firefox);
- **Advertisers**: companies with products and services to advertise to a targeted demographic (e.g., Squarespace);
- **Publishers**: websites that display ads to users viewing the webpage (e.g., `wired.com` or mobile applications);
- **Broker**: an ad-tech company (e.g., Google AdSense) responsible for matching users to ads, billing advertisers, and compensating publishers for user interactions.

The Broker, or ad network, can be thought of as the governing body of the advertising ecosystem. The Broker's primary goal is matching ads from an Advertiser to the users most likely to engage with them.

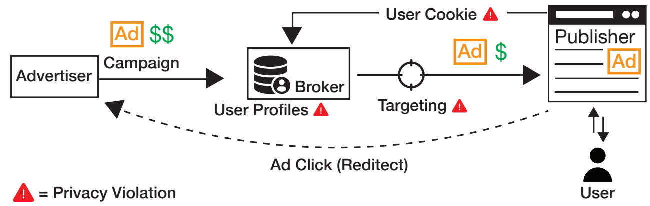

Figure 1: (non-private) targeted advertising pipeline

This is done via Publishers, companies that generate revenue by displaying ads to users visiting their webpage or app.

2.2 Advertising Pipeline

The advertising pipeline, shown in Figure 1, begins when an Advertiser creates a new ad campaign with the Broker, specifying a target audience and allocating a campaign budget. The target audience is specified using a *feature vector* which contains information about the contents of the ad, e.g., tags such as `mechanical keyboards` or `outdoor gear`. In addition to the feature vectors provided by Advertisers, the Broker possesses *user profiles* containing demographics and interests, e.g., `[woman, computers, high-income]`, about users. These profiles are constructed by tracking users on the web [98]. The Broker is responsible for *targeting* ads to users based on this profile through a similarity search algorithm [76].

When users are shown an ad on a publisher's webpage, a browser script generates an *impression* report for the Broker. This is by far the most common type of user engagement, followed by *clicks*, which are redirected through the Broker to measure engagement. With impression reports and click redirects, the Broker obtains necessary metrics for billing Advertisers, compensating Publishers, and updating targeting algorithms. Crucially, the Broker eliminates all fraudulent views and clicks that it deems to be generated by bots or malicious Publishers and clients to artificially skew perceived user interactions with ads [83].

2.3 Styles of Online Advertising

Online advertising consists of two main targeting strategies for matching ads to users: *contextual* and *behavioral* targeting.

Contextual advertising is done *independently of the user's profile*. Publishers display ads that are relevant to their own content, which is assumed to be relevant to visiting users based on their choice to access the Publisher's site or app.

Behavioral ad targeting matches a user to a set of ads based on the *user's profile*. Each profile consists of the user's browsing behavior or collection of apps installed on the user's device, obtained by the Broker through user tracking. Retargeting is a common type of behavioral advertising in which the Broker may preferentially display an Advertiser's ads to users who have previously interacted with that Advertiser, e.g., by visiting their website.

Remark 1. ***Real-Time Bidding*** *is a mechanism associated with both behavioral and contextual advertising. It occurs on-demand, at the time a user requests an ad, and auctions the available ad slot to the highest bidding Advertiser. Advertisers can take into account* *fine-grain information about the user, the webpage on which the ad will be shown, and even the time of day when making their bids.*

Each style of online advertising has tradeoffs between the relevance of ads and the invasiveness of the data required. Serving increasingly relevant ads, as in behavioral or retaregted ads, currently requires corresponding increases in user tracking granularity.

In the next section we describe how AdVeil supports the targeting of relevant ads *without* sacrificing user privacy.

3 THE ADVEIL SYSTEM

AdVeil is a privacy-preserving instantiation of the advertising pipeline. AdVeil is designed to be modular and extensible with each stage of the advertising pipeline implemented independently.

This allows AdVeil to provide a general-purpose solution for private advertising that maintains unlinkability guarantees for users. The Broker retains comparable targeting accuracy, performance, metrics, and fraud prevention to existing advertising pipelines *without* the ability to learn which ads a user sees. That is, the user's Personally Identifiable Information (PII) is kept locally on the user's client (i.e., the browser or device) and is never revealed in the clear.

3.1 Overview

AdVeil is designed around *targeting*, *delivery*, and *reporting*. These are the three stages of the advertising pipeline where personal user information is involved (see Figure 1). Our overview of AdVeil follows the steps described in Figure 2, which transforms the advertising pipeline of Figure 1 into a privacy-preserving system.

Targeting (steps 1a to 1c in Figure 2) assigns relevant ads to a user. The client holds a user profile (step 1a) described as a *feature vector*. Ads are likewise associated with a feature vector describing their targeting attributes. The targeting protocol reveals nothing to the Broker and outputs a selection of ad IDs to the user's client. In addition, for each targeted ID, the Broker provides a blind signature on a special anonymous token with an embedded private metadata bit [56] indicating whether the client is identified as a "human" or a "bot". The client later returns the unblinded token (step 1c) to the Broker in reporting, which the Broker uses to discard "bot" reports.

Delivery (steps 2a to 2c in Figure 2) retrieves a targeted ad to display to the user corresponding to an ad ID (step 2a). The choice of which targeted ID to request (from the set of IDs received in targeting) is performed by the client with the help of a local *selection function* (see Sections 5.2 and 9). The client then sends the ID to the Broker via the anonymizing proxy to hide its IP address. The output of the delivery protocol is a targeted ad selected by the Broker in real-time (step 2b) for the targeted ID. This mechanism provides support for real-time bidding and other on-demand delivery logic (see Section 5.3). As in Targeting, the Broker provides a blind signature on a fresh anonymous token with embedded *public* metadata bits [79], binding the token signature to the returned ad. Later, in reporting, the unblinded token (step 2c) ensures that the report is uniquely associated to a delivered ad by examining the metadata.

Reporting (steps 3a to 3c in Figure 2) provides interaction reports (e.g., impressions and clicks; step 3a) to the Broker. To do so privately, the client batches and sends all reports to the Broker at fixed time intervals through the anonymizing proxy. The client also attaches the signed tokens obtained in the targeting and delivery protocols. All reports that the Broker deems fraudulent based on these tokens are rejected. All remaining (i.e., valid) reports are used by the Broker for metrics and billing (step 3c).

3.2 Threat Model and Requirements

AdVeil has two independent security concerns: user privacy, which we define as *unlinkability* between users and their personal data, and *fraud prevention*, which we define as billing and reporting statistics that are resilient to botnets and false claims.

3.2.1 Personal Data. Concretely, we define **personal data** to be any information that is related, either directly or indirectly, to a user's targeting profile. This includes elements of the user's feature vector used for targeting, which may contain demographics and interests, and are thus *directly* related to the user's profile. It also includes any ads the user sees or interacts with. This is because, given the nature of ad targeting, ad features are *indirectly* related to the user's profile through the targeting algorithm. AdVeil never reveals any personal data directly, even anonymized or unlinked from the client's identity. The Broker may indirectly infer profile features from the billing statistics on ads, however, these remain unlinked from the identity of the corresponding user; see Section 6.

3.2.2 Adversarial Model. AdVeil assumes a rational [39, 44, 101] Broker that may try to link the identity of users to any of their personal data. A **rational** Broker is incentivized to provide correct functionality for its advertising ecosystem. The two assumptions we make about the Broker's behavior in order to provide user unlinkability in AdVeil are:

(1) the Broker does not deny service to honest users during ad targeting or delivery,
(2) the Broker includes all valid user reports when computing metrics on ad interactions.

The Broker has a direct financial incentive not to violate these two assumptions. If the Broker refuses to serve ads, then it cannot collect revenue from advertisers. Similarly, if some ad interactions from honest users are excluded from metrics, then the Broker violates the billing contract it has with Publishers and Advertisers.

3.2.3 User Unlinkability. User privacy in AdVeil is defined in the context of *unlinkability* between users and their personal data. AdVeil guarantees that any data, e.g., ad interactions, that users report on cannot be tied back to their identity. Users in AdVeil are only required to report ad interactions and, as such, their reports do **not** need to contain any information about their profile.

Unlinkability for users in AdVeil is defined on a per-epoch basis where each epoch is the time period from the start of the targeting protocol through the end of reporting. Epochs hide timing correlations between AdVeil's protocols and ensure that all users participating in a given epoch are part of an *anonymity set*. The Broker cannot determine which user within an anonymity set was responsible for generating a set of reports.[1]

[1]We discuss potential cross-epoch leakage resulting from intersection attacks [30, 31, 63, 65, 100] between the anonymity sets of different epochs in Section 6.1.3.

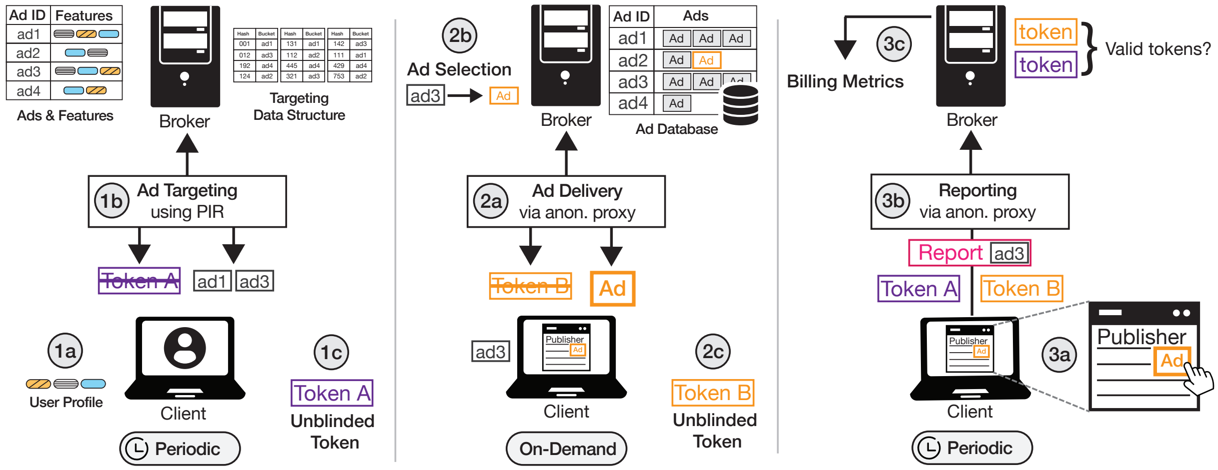

Figure 2: Overview of AdVeil. Description of the components is provided in Section 5.

We require that users make delivery and reporting requests via an anonymizing proxy such as Tor[2] [35]; see Section 4. This ensures that, while the Broker learns which ads are delivered and reported on, it does not learn the identity of the user who saw the ad.

3.2.4 Fraud Prevention. AdVeil expects users of the system to arbitrarily misbehave or to be impersonated by large scale bot networks. Users may try to avoid being shown ads at all or incur billing for ads that were never targeted or delivered.

AdVeil integrates a fraud-prevention mechanism using a combination of anonymous tokens [56, 79] that cleanly compose with existing bot detection mechanisms for fraud-prevention. Importantly, this is a *silent* method of fraud prevention. That is, while the Broker can recognize reports with invalid tokens, no other user can. This prevents bots from using AdVeil's tokens to learn which of their evasion strategies were successful.[3] In addition, AdVeil ensures that users only obtain tokens for the exact ads they were delivered and, as such, cannot generate valid reports on arbitrary ads.

In AdVeil, fraud-prevention guarantees that:

(1) the Broker can distinguish reports generated by a fraudulent request from those generated by an honest request,
(2) each valid report is included at most once in billing metrics,
(3) users can only generate valid reports for delivered ads.

We note that AdVeil does not make it easier for users to block ads. If users opt out of targeted advertising by not participating in AdVeil, then Publishers can fall back to *contextual* ads which do not require any user data or participation.

3.2.5 Non-goals. While AdVeil supports the return of arbitrary data during the reporting phase, it only requires reporting on the ads themselves for correct functionality. Determining what, if any, user data can be reported on privately is an orthogonal problem. AdVeil is compatible with any choice of data privacy mechanism such as differential privacy [36] or k-anonymity [87]. Similarly, while AdVeil guarantees that the ad identified in reports matches the ad delivered, we consider handling the integrity of any optional data submitted in reports to be out of scope.

Additionally, AdVeil is fully compatible with existing anti-tracking measures, however, it is not itself a method for thwarting web tracking. Instead, AdVeil serves primarily to disincentivize such tracking.

Finally, we cannot prevent users from intentionally disclosing PII to Advertisers and Publishers e.g., inputting their address to order an advertised product or logging into an account on a Publisher's website. AdVeil is neither intended to, nor capable of, preventing such behavior. As such, we see solutions for addressing this threat (e.g., through UI warnings and anonymous authentication) to be orthogonal research problems.

4 BUILDING BLOCKS

We design AdVeil around several standard cryptographic and data structure building blocks. Table 3 summarizes the building blocks and their use in AdVeil. Overall, we aim to make AdVeil general and modular, so as to fit a wide array of potential deployment scenarios.

For targeting, we use single-server PIR [7, 58] in conjunction with a standard similarity search data structure based on locality-sensitive hashing (LSH) [6, 40]. We emphasize that using a LSH-based data structure ensures that AdVeil is compatible with targeting techniques used in practice [76, 96].

We use an anonymizing proxy (such as Tor [35] or I2P [103]) for the purpose of fetching targeted ads from the Broker, which reveals which ad was retrieved but not which client retrieved it.

To guarantee integrity in both targeting and delivery, the Broker commits to all data structures using an authenticated dictionary (e.g., Merkle tree) which can be efficiently verified by the client.

For fraud prevention, we use anonymous tokens with a private metadata bit (PMBTokens) of Kreuter *et al.* [56] and subsequently extended by Silde and Strand [79] to support both *public and private* metadata. We use PMBTokens to enforce a "one-delivery; one report" policy for all reports and ensure that each report is uniquely associated with a prior delivery request.

[2]Any anonymizing proxy hiding the sender identity is sufficient; see Section 4.
[3]See https://github.com/WICG/trust-token-api#extension-metadata.

4.1 Private Information Retrieval

Private information retrieval (PIR) is a standard cryptographic technique for retrieving items from a remote database or data structure without revealing *which* item was retrieved [15, 25, 26]. PIR in the single-server setting [58] with a database of N items requires $O(N)$ work on the server and sublinear communication (in practice $O(\sqrt{N})$ communication [7]). In AdVeil, PIR is used by the client to privately query the targeting data structure.

Definition 1 (PIR [58] (informal))**.** *Let M be a set of instantiation-specific public parameters. For a fixed database $\mathcal{D}$ (or dictionary [25]), PIR consists of the following functionality:*

$\underline{\mathsf{PIR.Query}}(M, \mathsf{idx}) \to (s, Q)$. Query *takes as input public parameters M and index* idx. *Outputs secret state s and query Q.*

$\underline{\mathsf{PIR.Answer}}(\mathcal{D}, Q) \to A$. Answer *takes as input a database $\mathcal{D}$ and query Q. Outputs answer A.*

$\underline{\mathsf{PIR.Recover}}(s, A) \to a$. Recover *takes as input secret state s and query answer A. Outputs recovered database item a.*

where the functionality satisfies correctness *and* privacy.

Informally, correctness holds if the computed answer produces the item at index idx in the database $\mathcal{D}$ when fed through Recover. Privacy for the client holds if the (potentially malicious) server learns no information on the query. We refer to Kushilevitz and Ostrovsky [58] for formal functionality requirements.

4.2 Anonymous Tokens with Embedded Metadata

We make use of anonymous tokens with metadata bits (PMBTokens) [56, 79] for fraud-prevention in AdVeil. With PMBTokens, a verifier (e.g., the Broker) signs a blinded token generated by the prover (e.g., the client). During signing, the verifier can embed public and/or private metadata. Later, the unblinded token and the signature is redeemed by the verifier *without any linkability between signed and redeemed token*. Such tokens are a special form of blind signature, with the added property that the verifier can embed a metadata bit into the signature. We use both public and verifier-private metadata to achieve fraud prevention and rate-limiting in AdVeil.

Definition 2 (PMBToken [56, 79] (informal))**.** *A PMBToken scheme consists of efficient algorithms* KeyGen, TokenGen, Sign, Unblind *and* Redeem *with the following functionality:*

$\underline{\mathsf{PMBT.KeyGen}}(1^\lambda) \to (\mathsf{pk}, \mathsf{sk})$. KeyGen *outputs a new PMBToken public key* pk *and secret token signing key* sk.

$\underline{\mathsf{PMBT.TokenGen}}(\mathsf{pk}) \to (\tau, \bar{\tau}, r)$. TokenGen *outputs a new token τ, blinded token $\bar{\tau}$, and blinding factor r using the public key* pk.

$\underline{\mathsf{PMBT.Sign}}(\mathsf{sk}, \bar{\tau}, \mathsf{md}_{priv}, \mathsf{md}_{pub}) \to \bar{\sigma}$. Sign *takes as input the secret signing key, blinded token $\bar{\tau}$, private metadata md_{priv}, and public metadata md_{pub}. Outputs signature $\bar{\sigma}$ on the blinded token $\bar{\tau}$ with embedded metadata md_{priv} and md_{pub}.*

$\underline{\mathsf{PMBT.Unblind}}(\mathsf{pk}, \bar{\tau}, \bar{\sigma}, r) \to (\tau, \sigma)$. Unblind *takes as input the public key, blinded token $\bar{\tau}$, blind signature $\bar{\sigma}$, and blinding factor r. Outputs unblinded token τ and signature σ.*

$\underline{\mathsf{PMBT.Redeem}}(\mathsf{sk}, \tau, \sigma) \to (\mathsf{md}_{priv}, \mathsf{md}_{pub})$. Redeem *takes as input a secret key* sk, *token τ, and signature σ. Outputs private and public metadata* $(\mathsf{md}_{priv}, \mathsf{md}_{pub})$. *For notational convenience, we let* $(\mathsf{md}_{priv}, \mathsf{md}_{pub})$ *take on default values* $(0, 0)$ *if σ is invalid.*

The functionality must satisfy completeness, unlinkability, unforgeability, and metadata privacy (only for private metadata). A token may have no embedded metadata (public or private), in which case the metadata is denoted by the special symbol $\perp$.

Building Block	Purpose in AdVeil
LSH [6, 40]	Ad targeting based on user interests.
PIR [26, 58]	Private ad targeting with low bandwidth.
PMBToken [56, 79]	Fraud prevention in interaction reporting.
Authenticated Dictionary [70]	Integrity of targeting and delivery.
Anonymizing Proxy [35, 94, 103]	Ad delivery and private user interactions.

Table 3: Summary of building blocks and their use in AdVeil.

Loosely speaking, completeness and unlinkability state that a verifier always accepts valid tokens using Redeem but learns nothing beyond the metadata from the signed token redemption. Unforgeability states that a prover (i.e., the client) cannot forge a valid token without the secret key. Finally, metadata privacy (for private metadata only) requires the signed token reveal no information on the verifier-embedded metadata md_{priv} to the prover. For public metadata md_{pub}, the embedded metadata is readable by both the prover and the verifier.

Remark 2. *There is subtlety with respect to the private metadata, which in our case is a single bit, i.e.,* $\mathsf{md}_{priv} \in \{0, 1\}$. *Specifically,* md_{priv} *defines two sets: tokens that are* valid *and tokens that are* invalid. *Anyone can generate an invalid token (with* $\mathsf{md}_{priv} = 0$*). However, only the verifier can distinguish between these sets and generate tokens with* $\mathsf{md}_{priv} = 1$. *This ensures privacy of the embedded bit and, moreover, guarantees that the tokens define at most two groups. This observation becomes important when arguing unlinkability in Section 6, where* md_{priv} *partitions anonymity sets.*

4.3 Other tools

Authenticated Dictionaries [4, 13, 23, 70, 89] are a common technique for proving validity of a key lookup in a hash table relative to a short commitment. In AdVeil, we use authenticated dictionaries (e.g., a Merkle tree) to ensure lookup consistency across client requests. The two properties that are important for this work are that each proof (1) is of logarithmic (or constant [4, 13, 23, 89]) size relative to the size of the dictionary and commitment and (2) can be efficiently verified given only the commitment, lookup key, and proof string.

Anonymizing Proxies [35, 94, 103] serve as an intermediary, or *proxy*, for communications between a client and server to hide the relationship between sender and recipient of messages. AdVeil only requires *client anonymity* [10] from the proxy. That is, the identity (IP address) of the client should be hidden by the proxy. The Onion Router (Tor) [35], Invisible Internet Project (I2P) [103], and VPN0 [94] all provide this property. Of these, Tor is the most widely deployed and is bundled by default in the Brave browser [17] and made available through extensions for Firefox[4] and Chrome[5].

5 ADVEIL ARCHITECTURE

This section introduces our construction of AdVeil. We compose AdVeil from protocols that instantiate the stages of the advertising pipeline, using the tools of Section 4.

First, in Section 5.1, we detail non-private nearest neighbor search which forms the foundation of our private ad targeting protocol between the client and the Broker through PIR.

In Section 5.2, we describe the protocol constructions used for targeting, delivery, and reporting. In Section 5.3, we explain how the protocols combine to form a complete ecosystem for private targeted advertising.

5.1 Targeting and Nearest Neighbor Search

At the core of any targeting system is a nearest neighbor search data structure for matching ads to users [76, 104]. In practice, approximate hashing-based solutions are used to solve the neighbor search problem. Hashing-based data structures solve *approximate* nearest neighbor search (ANNS). Approximate solutions are necessary for efficiency given that exact solutions to the problem are believed to require brute-force search [50]. We note that ANNS is a *general* problem and can be applied to many definitions of similarity. Formally, ANNS is defined over a set of high-dimensional feature vectors and a query vector $\boldsymbol{q}$. For a fixed distance metric (e.g., Euclidean distance) ANNS returns the approximate nearest neighbor(s) to $\boldsymbol{q}$ (with respect to distance metric) from the set.

To see how ANNS defines a targeting system, consider a database $\mathcal{D}$ of N tuples of the form $(\mathsf{id}_i, \boldsymbol{v_i})$. Each tuple consists of an ad identifier id_i and a corresponding feature vector $\boldsymbol{v_i}$ describing the targeting attributes for the ad. To find the most relevant ad to a query $\boldsymbol{q}$ (where $\boldsymbol{q}$ is the user's profile vector), it suffices to find the nearest neighbor of $\boldsymbol{q}$, which we denote by $\boldsymbol{v_j}$, and output the corresponding ad ID id_j. This problem forms the basis for recommendation systems, including ad targeting [76, 96].

The data structure used for solving ANNS efficiently is based on locality-sensitive hashing (LSH) [6]. Common instantiations of LSH are MinHash [19] and SimHash [24]. An ANNS data structure is defined by two algorithms: Build and Query. Build takes as input the database $\mathcal{D}$, an LSH family $\mathcal{H}$, and a tunable parameter L. Build outputs ANNS data structure $\mathcal{S}$. Query takes as input the data structure $\mathcal{S}$ and a query vector $\boldsymbol{q}$. Query outputs the approximate nearest neighbors to $\boldsymbol{q}$ in $\mathcal{D}$ under the specified distance metric. Because our targeting protocol requires *privately* querying the ANNS data structure, we present a concrete instantiation of Build and Query as described in the seminal work of Gionis, Indyk, and Motwani [40]. In Section 5.2, we transform the Query function into a privacy-preserving protocol between the client and the Broker using PIR. We point to the survey of Andoni, Indyk, and Razenshteyn [6] for further details and discussion on parameter selection.

$\underline{\mathsf{ANNS.Build}(\mathcal{D}, \mathcal{H}, L)} \rightarrow \mathcal{S}$.

1: Define L hash tables $T_1, \ldots, T_L$ using LSH functions $\mathsf{h}_1, \ldots, \mathsf{h}_L$ sampled i.i.d. from LSH family $\mathcal{H}$; `// e.g.,` $\mathcal{H}$ `= MinHash or SimHash`

2: For each $\boldsymbol{v_i} \in \mathcal{D}$, compute $k_j \leftarrow \mathsf{h}_j(\boldsymbol{v_i})$, and append $(\mathsf{id}_i, \boldsymbol{v_i})$ to the bucket in hash table T_j with key k_j;

3: Output $\mathcal{S} = \{T_1, \ldots, T_L, \mathsf{h}_1, \ldots, \mathsf{h}_L\}$.

$\underline{\mathsf{ANNS.Query}(\mathcal{S}, \boldsymbol{q})} \rightarrow \mathsf{id}$.

1: Compute $k_j \leftarrow h_j(\boldsymbol{q})$ for $j \in [L]$;

2: Set $\mathcal{C} := B_1 \cup \cdots \cup B_L$ where each B_j is the bucket in hash table T_j with key k_j;

3: Output the id of the nearest neighbor to $\boldsymbol{q}$ in $\mathcal{C}$.

The key observation that we exploit for efficient (privacy preserving) targeting is that the ANNS data structure can be queried by the client using PIR (see Section 4.1), without revealing the sensitive query to the Broker. That is, the client can individually query each hash table using PIR and locally recover the nearest neighbor(s).

5.2 Protocols

AdVeil realizes each stage of the advertising pipeline via separate and modular protocols. We first describe the context in which the protocols are instantiated.

5.2.1 Setting. With the exception of Protocol 1, the protocols describe the steady state of AdVeil.[6] In this section, we begin by describing the background setting and starting assumptions.

One-time Trusted Setup. AdVeil (optionally) requires a trusted setup to instantiate authenticated dictionaries efficiently using a vector commitment [23]. While AdVeil is agnostic to the underlying authenticated dictionary instantiation, the most efficient (in terms of proof size) constructions require a trusted setup [42, 60].

Epochs. AdVeil divides time into discretized epochs. Aside from Protocol 1 which runs only once, all protocols run within the context of these epochs. The division of epochs is a deployment-specific parameter and only impacts the volume of users participating in each epoch. That is, epochs affect the frequency of ad targeting and reporting, but not *ad delivery* which remains on-demand.

Ad feature vectors. As a starting point, we assume that the Broker has a database of ads with associated feature vectors that describe the targeting attributes of each ad. How the feature vectors are obtained by the Broker is orthogonal to AdVeil. In practice, the feature vectors are provided by Advertisers or derived by the Broker through machine learning models. AdVeil is designed to be agnostic to the specifics of the targeting features in order to be compatible with a wide array of Broker strategies.

User profile. The user's profile feature vector is constructed locally by the client using a "profile building" function $\mathcal{F}_{\mathsf{profile}}$, provided as part of the public parameters generated in setup (Protocol 1). $\mathcal{F}_{\mathsf{profile}}$ maps user browsing history to a profile feature vector and can include information about websites they visit, searches performed, etc. [53, 57, 77, 96]. Ultimately, AdVeil leaves the final output of $\mathcal{F}_{\mathsf{profile}}$ up to the users themselves as we require they retain

[4]https://addons.mozilla.org/en-US/firefox/addon/tortm-browser-button/
[5]https://chrome.google.com/webstore/detail/onion-browser-button/fockhhgebmfjljjmjhbdgibcmofjbpca

[6]Protocol 1 can be re-run periodically to update targeting and ad database if needed.

Protocol 1: AdVeil Setup

Step 1 (Broker)

```
// 𝓛: feature vectors describing each ad
```

$\mathcal{S} \leftarrow \text{ANNS.Build}(\mathcal{L}, \text{params})$

$H_{\mathcal{S}} \leftarrow \text{AuthDict.Commit}(\mathcal{S})$ `// Commit to ANNS data structure`

$(\text{pk}, \text{sk}) \leftarrow \text{PMBT.KeyGen}(1^{\lambda})$ `// PMBToken keys`

Publish: public key pk, params, commitment $H_{\mathcal{S}}$, ad selection function $\mathcal{F}_{\text{ad}}$, profile building function $\mathcal{F}_{\text{profile}}$, and LSH functions $\mathsf{h}_1, \ldots, \mathsf{h}_L$ on a public bulletin board.

full control over the use of their data. Unlike in existing targeted advertising, this profile is not sent to or collected by the Broker.[7]

5.2.2 Setup (Protocol 1). The setup for the Broker involves publicly committing (e.g., through a public bulletin board [11, 27] or gossip network [59, 69, 71]) to all the parameters required for targeting, delivery, and reporting such that it cannot equivocate.

It begins with the Broker constructing the ANNS data structure for targeting. The Broker commits to the ANNS data structure hash tables and ad database (e.g., using a Merkle tree [70]), publicly releasing commitments to the authenticated dictionaries. The Broker then generates a PMBToken keypair and makes the public key available. For simplicity, we assume that all public parameters are consistent and accessible to all clients via a public bulletin board.

5.2.3 Targeting (Protocol 2). Targeting occurs at the beginning of each epoch (e.g., once a day) and runs between the client and the Broker. The client obtains the ID(s) of relevant ads while the Broker does not receive any output. We present retrieval of a single targeted ad. We note that Protocol 2 can be trivially extended to output the top-k ad IDs without significant overhead (see Section 5.1). The client uses the ID(s) of the targeted ads to fetch ad content in Protocol 3.

In addition to the ad ID(s), the client also receives the Broker's signature on a PMBToken (or k PMBTokens simultaneously). Each token contains an embedded bit, not observable to the client, that indicates whether the Broker believes the client to be a "bot" or "human". This bit is later used in Protocol 4 to reject bot reports.

To ensure integrity of targeting (i.e., that the Broker returns the correct set of targeted ads), the client checks the ID(s) it receives from a bucket in the ANNS hash table against the hash table commitment $H_{\mathcal{S}}$. To facilitate this, we assume the Broker concatenates the proof string for each bucket to the bucket contents. In this way, the client can retrieve the proof simultaneously with the bucket contents when privately querying the hash table through PIR.

5.2.4 Delivery (Protocol 3). Clients use Protocol 3 on-demand within epochs to retrieve ads corresponding to the ID(s) obtained from Protocol 2. The client fetches the corresponding ad from the Broker, via the anonymizing proxy. Importantly, the Broker has the ability to decide *which* ad to deliver based on the provided ID. We intentionally ensure that multiple ads can be associated with an ID so as to provide support for Broker-side selection logic (e.g., real-time bidding). We discuss PIR as an alternate delivery mechanism in Section 7, but note that it does not provide the required performance for on-demand delivery and ad selection logic.

[7]This mirrors other privacy-preserving targeting approaches [45, 73, 90, 96].

Protocol 2: Targeting

Step 1 (Client)

```
// p: user profile feature vector held by the client.
// h_i: LSH functions in the public parameters.
```

$k_i \leftarrow \mathsf{h}_i(\boldsymbol{p})$ for $i \in [L]$ `// Compute profile hash keys`

$Q_i \leftarrow \text{PIR.Query}(M, k_i)$ for $i \in [L]$ `// Query for bucket`

$(\tau_t, \cancel{\tau_T}, r) \leftarrow \text{PMBT.GenToken}(\text{pk})$ `// New reporting token`

Send queries $Q_1 \ldots Q_L$ and $\cancel{\tau_T}$ to the Broker.

Step 2 (Broker)

```
// 𝒮: ANNS with (authenticated) hash tables {T_1, ..., T_L}
```

$A_i \leftarrow \text{PIR.Answer}(T_i, Q_i)$ for $i \in [L]$ `// Compute query answer`

$\text{md}_{priv} \leftarrow \begin{cases} 0 & \text{if client identified as } \textbf{bot} \text{ // Invalid token} \\ 1 & \text{if client identified as } \textbf{human} \text{ // Valid token} \end{cases}$

$\cancel{\sigma_T} \leftarrow \text{PMBT.Sign}(\text{sk}, \cancel{\tau_T}, \text{md}_{priv}, \perp)$

Send $A_1, \ldots, A_L$ and $\cancel{\sigma_T}$ to the client.

Step 3 (Client)

$(B_1 \| \pi_1, \ldots, B_L \| \pi_L) \leftarrow \text{PIR.Recover}(A_1, \ldots, A_L)$

If there exists i such that $\text{AuthDict.Verify}(H_{\mathcal{S}}, B_i, \pi_i) = \text{no}$, abort.

$\text{id} \leftarrow$ nearest neighbor to $\boldsymbol{p}$ in $\mathcal{C}$ where $\mathcal{C} := B_1 \cup \cdots \cup B_L$.

$(\tau_t, \sigma_t) \leftarrow \text{PMBT.Unblind}(\text{pk}, \cancel{\tau_T}, \cancel{\sigma_T}, r)$ `// Unblind signature`

Output $(\text{id}, \tau_t, \sigma_t)$

Protocol 3: Delivery

Step 1 (Client)

```
// id : targeted ID obtained from Protocol 2.
```

$(\tau_d, \cancel{\tau_d}, r) \leftarrow \text{PMBT.GenToken}(\text{pk})$ `// New delivery token`

Send $(\text{id}, \cancel{\tau_d})$ to the Broker, via the **proxy**.

Step 2 (Broker)

```
// 𝒟: database of ads
```

$\text{ad} \leftarrow \text{SelectAd}(\mathcal{D}, \text{id})$. `// Broker's ad selection logic`

$\cancel{\sigma_d} \leftarrow \text{PMBT.Sign}(\text{sk}, \cancel{\tau_d}, \perp, \text{id})$ `// Sign with public metadata id`

Send $(\text{ad}, \cancel{\sigma_d})$ to the client, via the **proxy**.

Step 3 (Client)

$(\tau_d, \sigma_d) \leftarrow \text{PMBT.Unblind}(\text{pk}, \cancel{\tau_d}, \cancel{\sigma_d}, r)$ `// Unblind signature`

If $\text{PMBT.Verify}(\text{pk}, \sigma_d, \text{id}) = \text{no}$, abort. `// Invalid public metadata`

Output $(\text{ad}, \tau_d, \sigma_d)$.

5.2.5 Reporting (Protocol 4). Reporting occurs at the end of an epoch. When the user interacts with a retrieved ad through a Publisher's webpage (e.g., by viewing the ad or clicking on it), an impression report is generated and stored by the client. At the end of the epoch, the client sends each report to the Broker, via the anonymizing proxy, as specified in Protocol 4. The client includes the signed tokens obtained in Protocols 2 & 3. Upon receiving a report, the Broker verifies the tokens, rejecting all reports flagged as "bot" in targeting. All valid, non-duplicate reports are accepted.

Protocol 4: Reporting

Step 1 (Client)

```
// report: report payload.
// (τ_t, σ_t): PMBToken signature obtained from Protocol 2.
// (τ_d, σ_d): PMBToken signature obtained from Protocol 3.
```

1: Send $(\text{report}, \tau_t, \sigma_t, \tau_d, \sigma_d)$ to the Broker, via the **proxy**.

Step 2 (Broker)

```
// sk: secret key of the PMBToken scheme.
// W: list of all redeemed tokens.
```

1: $(\mathsf{md}_{priv}, \bot) \leftarrow \mathsf{PMBT.Redeem}(\mathsf{sk}, \tau_t, \sigma_t)$ `// Private bit`
$(\bot, \mathsf{id}') \leftarrow \mathsf{PMBT.Redeem}(\mathsf{sk}, \tau_d, \sigma_d)$ `// Public metadata`

2: If $\mathsf{md}_{priv} = 0$: reject. `// Bot token`

3: If $\mathsf{id}' = \bot$: reject. `// Invalid delivery token`

4: If $\tau_d \in \mathcal{W}$: reject. `// Duplicate delivery token`

5: Else, add τ_d to $\mathcal{W}$ and accept.

5.3 The AdVeil Ecosystem

AdVeil focuses on *general* targeted advertising, with support for multiple targeting and metrics strategies.

Targeting Strategies. By giving the Broker control over the choice of ANNS data structure and real-time ad selection logic, AdVeil does not restrict the Broker to a specific targeting algorithm while still ensuring privacy for users. Types of targeting that can be supported in AdVeil include:

- **Contextual** targeting in AdVeil can bypass users entirely by having the Publisher fill their role in the protocols. Users will still see and be able to interact with ads, but will have no involvement in any other aspect of the pipeline.
- **Behavioral** targeting in AdVeil is achieved through the profile building function $\mathcal{F}_{\text{profile}}$ which is able to locally observe and record information about users' browsing habits forming a profile feature vector for the user. Protocol 2 is then run periodically to retrieve a new set of targeted ads based on the user's profile. Retargeting, or preferentially displaying ads to users who have had prior interactions with the Advertiser, is one example of behavioral targeting that AdVeil can support in this manner.

As mentioned in Section 5.2, the only requirement for targeting strategies is that they must be fully local. Users must not make any requests other than those specified in Protocol 2, Protocol 3, and Protocol 4 as a consequence of any targeting strategy.

Real-Time Bidding in AdVeil is supported automatically by giving the Broker full control over which ad to deliver to the client based on the ad ID requested. The Broker can, in real time, select from a set of ads to respond with in Protocol 3 for a given ad ID. That is, while the targeted ID is fixed, the associated ad content need not be. Because the delivery and reporting requests are unlinkable from the targeting request, delivering different ads for the same ad ID does not compromise unlinkability, see Section 6.

Measurement Strategies. Reports in AdVeil are *individual* without being *linkable*. The Broker learns precisely which ads were seen and how often, without ever learning who saw them. This allows the Broker to support a variety of measurement strategies including:

- **Impression** reports generated when users *see* a displayed ad.
- **Click** reports generated when users *click* on a displayed ad.
- **Conversion** reports generated when a user engages with the Advertiser after clicking on an ad. The user's local client is capable of observing when an ad click generates a conversion event and creates the resulting report.[8]

All reports in AdVeil are assumed to contain, at minimum, the ad delivered and interaction type (e.g., Impression or Click). There is a large body of *other* information the Broker may wish to receive as part of reporting. AdVeil supports reporting on arbitrary information but, as discussed in Section 3.2, is not intended to provide *data* privacy guarantees. While it is always possible to report on user data privately using differential privacy [36] or k-anonymity [87], such techniques are deployment specific and orthogonal to AdVeil's primary goal of providing unlinkability between users and reports.

6 SECURITY ARGUMENTS

In this section we analyze the security of AdVeil when instantiated using the protocols of Section 5.2. We frame our analysis in terms of *user unlinkability* and *fraud prevention*, the two requirements outlined in the threat model of Section 3.2.

6.1 User Unlinkability

We recall that the security guarantee of AdVeil, as established in Section 3.2, is *unlinkability* between users and their personal data, including which ads they see and interact with.

We show in Claims 1 and 2 that, individually, none of the protocols are linkable to personal user information. Therefore, the crux of the unlinkability argument lies in analyzing the combination of targeting (where the Broker learns the client's identity; required for effective fraud-prevention) and delivery/reporting (where the client's identity is hidden by the anonymizing proxy). Concretely, we must ensure that the Broker cannot deviate in targeting to link the client in delivery or reporting, without also compromising its own goals (violating the *rationality* assumption; see Section 3.2).

Theorem 1 (Unlinkability). *Fix a reporting epoch. The set of recovered reports through Protocol 4 in the epoch is unlinkable from the set of clients that submitted them as well as prior executions of Protocol 2 and Protocol 3, conditioned on:*

(1) the user and client do not explicitly or implicitly leak personally identifying information to any party,

(2) the Broker is rational in the sense that it does not deny service to honest users or self-sabotage the fraud-prevention mechanism,

(3) the privacy of PIR [58], soundness of the authenticated dictionary [70], unlinkability of PMBTokens [56, 79], and client anonymity property of the anonymizing proxy [10], all hold under their respective assumptions.

6.1.1 *Privacy of individual protocols.*

We first show that the targeting, delivery, and reporting protocols, *individually* provide unlinkability from user data.

Claim 1. *Protocol 2 (targeting) reveals the identity of the client but no other information to the Broker, conditioned on the privacy requirement of the PIR scheme.*

[8] This is similar to how conversions are tracked in the Safari browser today [99].

Proof. The claim follows almost immediately from protocol inspection. Specifically, in Protocol 2 (targeting), the client only interacts with the Broker through a series of PIR queries which, by definition, hide the user's query (i.e., profile) from the Broker. The PMBToken issued in Protocol 2 is generated by the client independently of all user data and hence reveals no information on its own. ■

Claim 2. *Protocol 3 (delivery) and Protocol 4 (reporting) reveal the ad delivered and report contents to the Broker, but not the identity of the client, conditioned on the client anonymity property of the anonymizing proxy [10].*

Proof. In Protocols 3 & 4, the client interacts with the Broker through anonymizing proxy, which reveals the targeted ad ID and which ad was served (resp. the contents of the report) but not which client it was served to (resp. which client submitted the report). This follows directly from the client anonymity property of the anonymizing proxy [10]. The PMBTokens generated and sent by the client in both the delivery and the reporting protocols contain no identifiable information as they are computed independently of the user data.

■

6.1.2 Proof of Theorem 1. The unlinkability argument hinges on showing that a malicious (but rational) Broker cannot deviate in the targeting (Protocol 2) to link the client in delivery (Protocol 3), or reporting (Protocol 4). This is shown in Lemma 1.

Lemma 1. *Fix a reporting epoch. The Broker cannot rationally deviate from Protocol 2 to link the client to a delivery request (Protocol 3) or report submitted through Protocol 4, assuming the soundness of the authenticated dictionary [70] and the unlinkability property of PMBTokens [56, 79].*

Proof. At a high level, Lemma 1 follows from the Broker committing to the ANNS dictionaries in Protocol 1 (setup) and the total number of ads, coupled with the unlinkability property of PMBTokens. More formally, we must individually examine interactions taken by the client and the Broker in an epoch.

First, because the ANNS dictionaries are committed to in Protocol 1, a malicious Broker cannot change the ANNS data structure between client requests. As such, the contents of the targeting data structures must be *consistent* across all clients and, consequently, the contents of all clients' reports must also be consistent with the fixed targeting data structures. This is guaranteed by the authenticated dictionary proof returned to the clients with their PIR queries. If the Broker is capable of answering the PIR query with a valid proof (w.r.t. the commitment H_S) for a bucket value that is *not* in the dictionary, with better than negligible probability in a security parameter, then the Broker is also capable of breaking the soundness property of the authenticated dictionary with non-negligible probability [70].

Second, we examine the PMBToken (and the embedded private metadata) as a vector for linking a client to a report. We show that exploiting PMBTokens for linking clients to reports, while possible, is monetarily disincentivized and hence falls under irrational behavior. Specifically, because each token issued in Protocol 2 is set up to reveal *one* bit of metadata to the Broker for fraud-prevention purposes, the Broker can only partition the anonymity set into two groups: clients with valid tokens and clients with invalid tokens (recall Remark 2). This division is necessary for any scheme that allows the Broker to silently tag fraudulent requests for later identification. However, all tokens within these two sets are unlinkable from other tokens in their respective set by the properties of PMBTokens [56]; hence the Broker cannot link a valid (resp. invalid) token to a prior targeting request.

A malicious Broker can still split a single client into their own group (by issuing only one valid token), allowing it to directly link their identity to the contents of their report. However, this sabotages fraud-prevention, as the Broker must group all other users **and** all bots into the other set (invalid tokens), losing its ability to distinguish fraudulent reports and causing it to violate its billing arrangement with Advertisers and Publishers. This strategy would result in significant monetary losses for the Broker, as it sabotages correctness for its advertising ecosystem. As such, a *rational* Broker will use the metadata bit in the tokens only to tag fraudulent requests.[9] ■

As a consequence of Lemma 1, the Broker cannot deviate from targeting without either compromising fraud-prevention or denying service to a client (i.e., causing the client to abort). Such behavior, however, is inherently irrational as it either 1) prevents users from being shown ads or 2) sabotages fraud-prevention resulting in potentially inaccurate billing metrics. As a result, a rational (monetarily-driven) Broker is incentivized to follow protocol ensuring the unlinkability property is satisfied between clients and generated reports.

6.1.3 Cross-Epoch Unlinkability. If only a subset of all users participate in each epoch, then AdVeil cannot provide full unlinkability across epochs due to intersection attacks [30, 31, 63, 65, 100]. In this case, intersection attacks become possible because users are correlated with the ads they see. Thus, over time, the Broker can infer the relationship between a user and the ads they are shown by observing the *intersection* of epochs in which the user participated, *even though the unlinkability property is satisfied within each epoch.* That is, ads that appear most frequently alongside a certain user are likely to be the ads that user reported on. This leakage is small, but not resolvable without holding the set of either users or ads constant across all epochs – neither of which is reasonable for an internet-scale system.

6.1.4 Unlinkability of Features. AdVeil guarantees that the Broker cannot link any user to any of their personal data, but does not make guarantees about any associations between the data elements themselves. That is, the Broker may learn relationships between different *features*, e.g. that "`mechanical keyboards`" and "`programming`" are two distinct features that commonly occur together in user profiles.

More generally, the Broker can determine whether there exists *some* user who has a specific set of n features by generating a "tagged" ad ID that is targeted to exactly that set of features and no others. The Broker can then observe whether a report corresponding to the tagged ad ID appears during the reporting phase. If so, there exists at least one user that has the tagged set of n features in their profile. This does not reveal information about *which* user or users are involved,

[9] We note that the private metadata bit is only necessary to support *covert* bot detection. In situations where covert detection is not required, it is possible to dispense of the private metadata by e.g., using vanilla anonymous tokens such as Privacy Pass [34].

but allows the Broker to learn feature clusters, or sets of features that commonly appear together in user profiles.

6.2 Fraud Prevention

Recall that the requirements for fraud prevention (detailed in Section 3.2), are that 1) the Broker is capable of flagging all targeting requests detected as coming from bot clients and 2) no duplicate or bot reports are recovered through Protocol 4.

THEOREM 2 (FRAUD PREVENTION). *Assuming the unforgeability property of PMBtokens [56, 79], each report recovered by the Broker through Protocol 4 is:*

(1) submitted by a client given a valid token signature in Protocol 2,
(2) unique in the set of all reports submitted across all epochs,
(3) associated with an ad that was delivered in Protocol 3,

PROOF. The association of each report to a unique prior execution of Protocol 3 is guaranteed by the unforgeability property of the PMBTokens; no client can forge a valid token signature without knowledge of the secret signing key held by the Broker [56]. Report uniqueness is likewise guaranteed by the PMBToken unforgeability property and the dictionary $\mathcal{W}$ of all redeemed tokens maintained in Protocol 4 (see Privacy Pass [34] for more details and optimizations). Finally, the PMBToken unforgeability property guarantees that no client can forge a valid signature on a token. To expand on this final point, only clients marked as "human" in Protocol 2 are given valid signatures on the token. All other clients are given "invalid" signatures (which are indistinguishable from valid signatures to the client); in essence, valid/invalid signatures encode a private metadata bit [56] (see Remark 2). In Protocol 3, the Broker issues a signature with *public* metadata containing the delivered ad ID. Due to the unforgeability property [79], this fixes the ad ID, which in turn means that the client cannot report on a different ad from the one requested through Protocol 3. ■

6.3 Correctness and Efficiency

In this section, we briefly argue correctness and asymptotic efficiency of the protocols composing ADVEIL.

6.3.1 Correctness. Correctness of ADVEIL follows immediately from the correctness of the underlying building blocks. Accuracy of targeting is inherited from the correctness of the similarity-search data structure of Section 5.1 and PIR [26]. Correctness of delivery follows from correctness of the anonymizing proxy routing traffic between the client and the Broker. For reporting, we note that the client submits the report which, at minimum contains the ad ID embedded in the token signed in Protocol 3 (delivery). By the correctness property of PMBTokens (signed in targeting and delivery), only valid reports are accepted by the Broker. The correctness of the recovered metrics thus follows from the correctness of the targeting and delivery protocols.

6.3.2 Asymptotic Efficiency. We provide a brief analysis of *asymptotic* efficiency of ADVEIL, in terms of the efficiency of the:

(1) ANNS data structure used in targeting,
(2) PIR scheme used to query the data structure,
(3) anonymizing proxy used in delivery and reporting,
(4) reporting token redemption and storage overheads.

Only the first three factors impact the efficiency for the end user. Asymptotically, the guarantees of the similarity search data structure [40], and the underlying PIR scheme [7, 58], result in communication of $O(N^{\epsilon})$ for $\epsilon > 0$. In practice, $\epsilon \approx \frac{1}{2}$ [6, 7]. If the anonymizing proxy is instantiated using Tor [35], then the overhead on the client is $O(1)$ in delivery and reporting. The Broker's work for targeting ads is $O(N)$ due to lower-bounds in PIR [26] and $O(1)$ for delivery and reporting (dictionary lookup). The work and space imposed on the Broker in reporting is $O(R)$ for an epoch containing R submitted reports in total (note: in practice PMBToken key rotation[10] can be used to prevent storing reporting tokens across different epochs for de-duplication purposes [79]).

7 EVALUATION

We implement a prototype of ADVEIL and measure its end-to-end performance in a networked deployment. We evaluate the computational overhead of the Broker when serving client requests and the end-to-end latency of targeting and delivery on the client. We note that we do not compare ADVEIL *quantitatively* with existing work in privacy-preserving advertising. Prior proposals are either of a theoretical/qualitative nature [52, 97], only solve one aspect of the pipeline under different assumptions [43, 47, 91], or have incomparable approaches [45, 47, 90]. We instead *qualitatively* compare all these systems with ADVEIL in Section 8.

7.1 Setup

Implementation. We implement ADVEIL in Go (v1.13) and C++17. Our implementation is open source and available at http://adveil.com/code. We use the open-source Microsoft SealPIR library [7] to instantiate single-server PIR. Our implementation of PMBTokens is partially based on Cloudflare's Privacy Pass code [34]. We instantiate authenticated dictionaries using a vector commitment [60], which require a one-time trusted setup [23] but have very succinct opening proof sizes of 48 B [23, 42, 60]. Concretely, short proofs improve the performance of targeting where the user retrieves the proof in conjunction with the value from the dictionary via PIR.

Environment. We deploy ADVEIL on Amazon Elastic Cloud Compute (EC2) server instance for the Broker and a MacBook Pro for the client. The Broker's server runs on `c5a.8xlarge` VM (32 vCPUs @ 3.6GHz; 64 GB RAM) with an hourly cost of \$0.525 as of June 2021.[11] We run the client on a MacBook Pro (8 CPUs; 32 GB RAM). We measure a throughput of 200 Mbit/s download and 5 Mbit/s upload using the `speedtest-cli` tool.[12]

Parameters. The tunable parameters in ADVEIL include the number of ads, the size (in KiloBytes) of each ad, and the number of hash tables used in the ANNS data structure (see Section 5.1). The total number of ads and ad size has a direct impact on network bandwidth usage and server processing time due to PIR (see Section 4.1). Likewise, for the ANNS data structure, more hash tables improve accuracy, but also increase the total number of parallel PIR queries required (see Protocol 2). In our runtime experiments, we vary the number of hash tables from $L = 5$ to $L = 30$. We report the accuracy

[10] See https://engineering.fb.com/2021/04/16/security/dit/
[11] https://aws.amazon.com/ec2/spot/pricing/
[12] https://www.speedtest.net/apps/cli

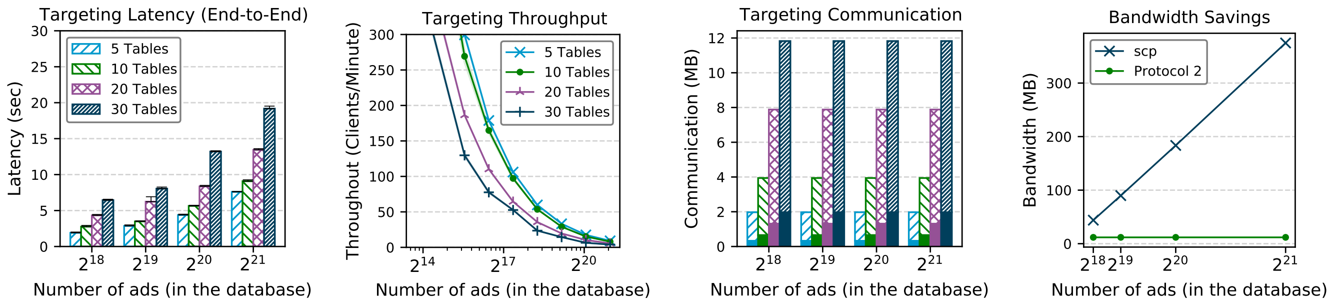

(a) Client end-to-end latency (seconds) for a targeting query. Latency is primarily impacted by Broker computation time (computing query answers), but also affected by ping time and client network bandwidth.

(b) Broker targeting throughput (clients per minute) with 32 cores. Throughput is only impacted by PIR query answer computation which is linear in the number of ads in the database.

(c) Client-Broker communication for targeting (PIR query and response size). Bottom solid areas of each bar indicate fraction of upload communication.

(d) Communication comparison between "naive" PIR (retrieving entire ANNS data structure (`scp`)) and Protocol 2, which uses SealPIR queries. We fix $L = 30$.

Figure 4: Evaluation of ad targeting in AdVeil. We report end-to-end latency as measured on the client machine. Communication is measured as the network overhead between the client and Broker. Throughput measures the raw computational processing throughput of the Broker's server (parallelized across 32 cores) to answer client targeting requests. Shaded regions and error bars represent a 95% confidence (occasionally invisible).

of the ANNS protocol while varying the parameter L in Figure 6. We make each hash table in the ANNS data structure have N keys (i.e., total number of ad IDs) and cap the size of each bucket in the ANNS to one ID.

Datasets. We evaluate AdVeil using synthetic data. Real-world ad features are proprietary [96] and we were not able to find suitable data for our evaluation. We stress, however, that PIR (which is used to query the ANNS hash tables) is not impacted by the underlying data distribution. Moreover, AdVeil does not impact *accuracy* of targeting and hence is also agnostic to the underlying features.

In contrast, the number of hash tables in the ANNS *does* impact accuracy (which directly impacts AdVeil's performance by requiring more PIR queries). To gain a sense of how many hash tables are required in a deployed setting, we follow the synthetic data generation and evaluation methodology performed by Datar *et al.* [32] for evaluating ANNS queries over Euclidean distance in a worst-case manner. Each ad feature vector is randomly generated by sampling a 100-dimensional vector with random coordinates in the range $[-255, 255]$ (each coordinate is one byte). A client profile feature vector is then randomly "planted" within a fixed distance radius from a randomly chosen ad feature vector. Because LSH performs best on *clustered* data (as is often the case with non-synthetic data), generating the dataset in this way results in *worst-case* recall performance for ANNS [6]. This setup is used to evaluate *accuracy* of approximate nearest neighbor search on specific LSH parameters[13] as a function of the number of hash tables to gain a crude estimate for the number of tables required in practice.

Ad Sizes. Common online banner ad sizes[14] (which account for 89% of ads on some platforms [64, 68]) are typically under 150 kB in size. Video ads are typically under 4 MB in size [68]. Ad size only impacts delivery as all other protocols operate on fixed-size ad IDs.

Methodology. Unless otherwise stated, we run each experiment 10 times and report mean and 95% confidence interval over the trials. We parallelize all computation on the servers when possible.

7.2 Results

In this section we describe our evaluation results for targeting, delivery, and metric recovery in terms of processing time, latency, and communication. We also report the impact of changing the number of LSH hash tables has on targeting accuracy.

7.2.1 Client latency. We first evaluate client end-to-end latency in wall-clock time for ad targeting and delivery.

Targeting (Protocol 2) Latency. Figure 4a shows the impact that the number of ads in the dataset and number of ANNS hash tables has on client latency. Latency ranges from several seconds on smaller sets of ads ($\approx 130,000$ ads) to 20 seconds on larger sets ($\approx 2,000,000$ ads). Network throughput accounted for 1 to 7 seconds of end-to-end latency on the client (which has throughput comparable to a home WiFi network).

Delivery (Protocol 3) Latency. Figure 5 shows the relative performance of Tor vs. single and two server PIR for delivering ads. We take the mean latency for downloading a 50 kB and 1 MB file over the Tor network from Tor metrics [62] data between February and May 2021. Delivery latency for a 50 kB ad (e.g., a small image) through Tor is approximately one second while latency for a 1 MB ad (e.g., 5 second 480p video[15]) is approximately three seconds.

To illustrate the impracticality of using PIR for delivering ads, we evaluate SealPIR over 500 B to 1 kB ads while varying the total number of ads in the database. Additionally, we consider two-server PIR, which is more concretely efficient in practice [15, 95]. Our results show that both single-server and two-server[16] PIR impose a high overhead for delivery, even when only considering small (up to

[13] Many factors are at play when determining optimal LSH parameters; see [81] for details on parameter optimization for ANNS in practice.

[14] https://support.google.com/adsense/answer/6002621

[15] See https://www.animotica.com/blog/full-guide-what-is-video-bitrate-and-why-does-it-matter/ for video bitrates.

[16] The two servers are required to be non-colluding in this setting.

1 kB) ads. As such, we choose to use Tor for ad delivery in AdVeil. We stress that *targeting* requires use of PIR as replacing it with Tor (or any other anonymizing proxy) would require user data to be sent to the Broker during targeting, which would directly reveal PII.

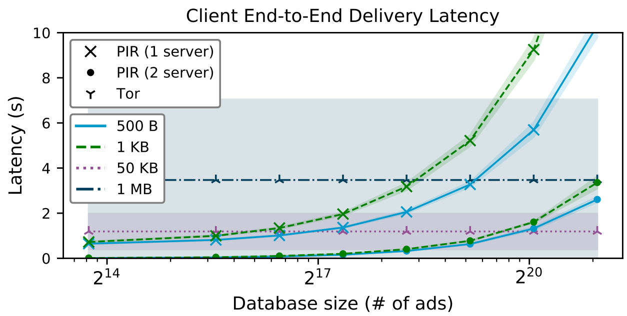

Figure 5: Client end-to-end ad delivery latency (seconds) using PIR, two-server PIR, and Tor. Latency for PIR queries is proportional to the size of the ad database (number of total ads and their size). Latency for Tor is only dependent on ad size (independent of the total number of ads). Shaded region represents a 95% confidence interval (note that Tor has high variance in latency).

7.2.2 Communication overhead. We report the total communication (in MB) exchanged between the Broker and the client when targeting ads in Figure 4c. Communication depends on database size and number of tables. The total communication remained under 12 MB. We find that the majority of communication is from the PIR query response since the PIR queries themselves are of constant size with respect to the database size in our evaluation (due to SealPIR query compression [7]). We believe this communication to be reasonable for the average internet client, especially considering targeting happens once per epoch (e.g., once a day) and is "download-heavy", which aligns well with real-world networks [82]. In Figure 4d we contrast the bandwidth usage of SealPIR with "naive PIR" where the entire ANNS data structure is sent to the client.

7.2.3 Targeting throughput. We report the targeting throughput (in terms of clients per minute) of the Broker in Figure 4b. Throughput is limited by the computational overhead of processing the PIR queries over the ANNS hash tables. With 30 hash tables and 2^{20} ads, targeting throughput was approximately six clients per minute. We note that the throughput is massively parallelizable; increasing linearly with the Broker's computational capacity. If we assume that targeting is required once per week and 100 million active users in the ecosystem (with anonymity sets of 2^{20} unique ad IDs), the Broker would need $\approx$ 1,600 servers to target ads on a weekly basis.

7.2.4 Targeting Accuracy. We compare targeting accuracy in AdVeil to that of contextual advertising as users' features become increasingly distant from those of the website. ANNS accuracy is typically measured through *recall*: the fraction of neighbors found over the total number of queries performed. While AdVeil is agnostic to the ANNS parameters (such as the LSH family used) we nonetheless measure the accuracy of nearest neighbor queries as a function of the number of hash tables used in the ANNS data structure. We first fine-tune the ANNS parameters to achieve recall accuracy of over 90% with $L = 30$ hash tables and report the drop in recall as we decrease the number of hash tables. We report recall accuracy across different values of L in Figure 6, evaluated over 10,000 independent queries while varying the number of hash tables in the ANNS data structure. We find diminishing returns in recall improvement when the number of hash tables increases over 20; suggesting that on real-world data (not worst-case as evaluated here), 20 hash tables in the ANNS may be sufficient for accurate targeting.

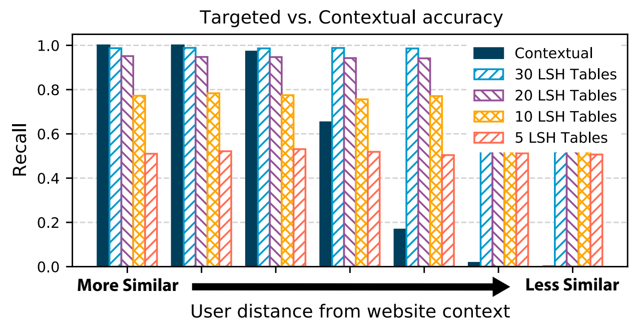

Figure 6: Recall accuracy comparison between targeted and contextual ads as a function of the user's profile distance from the website's profile. Increasing the number of hash tables has diminishing returns on recall accuracy for targeted ads. As user profiles grows less similar (i.e., more distant) from the website profile, contextual advertisement become less relevant. Targeted advertisement relevancy is independent of the user's similarity to the website context.

7.2.5 Reporting. The computation on the Broker per submitted reports is light and consists of redeeming the two tokens attached to each report. In total, redeeming the tokens for a single report requires under 300 μs of processing time on one core (see Table 7). Thus, a single 32 core server can process upwards of 90,000 reports per second when parallelized.

7.2.6 Microbenchmarks. We execute a series of microbenchmarks to analyze the impact of the building-blocks used in AdVeil. Specifically, we measure the client processing time for PIR queries and PMBTokens. Generating the PIR queries requires under one millisecond. Decryption of the PIR query response takes approximately two milliseconds. All client-side PMBToken processing (generating and unblinding) is under 300 μs.

Server-side processing of PMBTokens requires under 300 μs for all operations. PMBTokens with 1-bit of private metadata (compared to public metadata) are more efficient as they do require hashing of the metadata and zero-knowledge proofs during the signing phase [56].

7.2.7 Storage Overhead. We report the storage overhead in Table 8. The client stores the user's profile feature vector locally which requires d bytes (assuming one byte per profile vector coordinate). Additionally, the client stores the PIR key pair and public parameters published by the Broker in Protocol 1. The Broker stores all redeemed tokens to prevent duplicate reports. Additionally, we assume the Broker stores the PIR public keys of each client to avoid having the client to repeatedly send them with each PIR query.

Microbenchmarks			
Client			
879 μs	2007 μs	219 μs	272 μs
PIR.Query	PIR.Recover	PMBT.GenToken	PMBT.Unblind
Broker			
51 μs	94 μs	172 μs	105 μs
PMBT.Sign (private metadata)	PMBT.Redeem (private metadata)	PMBT.Sign (public metadata)	PMBT.Redeem (public metadata)

Table 7: Microbenchmarks (in microseconds) for SealPIR (query and answer recovery) and PMBTokens (with private and public metadata). PMBTokens with public metadata require additional processing due to hashing and non-interactive zero-knowledge proofs.

Storage Overhead				
Client			**Broker**	
d (B)	4 (MB)	200 (KB)	$32R$ (B)	$4C$ (MB)
Profile vector	SealPIR key	Public params	Report dict.	Client keys

Table 8: Storage overhead. C **is the total number of clients in the system.** R **is the total number of reports submitted in an epoch.** $d = 100$ **in our evaluation.**

7.2.8 Operational Costs. We estimate the operational costs of running AdVeil. Our estimate focuses on an ad database and reports-per-epoch size of 1M ads. We assume users recovers 10 ads simultaneously in Protocol 2 using 30 LSH tables (see Section 5.1 for how ANNS trivially extends to k-nearest neighbor search). Additionally, we use the AWS costs from our own evaluation – \$0.525/hr for the Broker's machine on AWS. Given this setup, the processing time of the Broker is approximately 10s per execution of the targeting protocol (Protocol 2) [17]. The total costs to target, and recover reports for, $k = 1$ ads is computed as in Equation (1) and equals 0.15¢.

$$\underbrace{(\$0.525/\textsc{hour}) \cdot 10\text{ s}}_{\text{cost of targeting } k \text{ ads}} + \underbrace{k \cdot (\$0.525/\textsc{hour} \cdot 0.5\text{ ms})}_{\text{PMBTokens cost for } k \text{ ads}} \quad (1)$$

Compared to average revenue generated by an ad impression of approximately 0.20¢ [84], we get that the expected revenue is roughly 33% the cost of serving the ads (on non-enterprise machines). Amortized over $k = 10$ ads targeted simultaneously, the expected revenue is over 10× the cost of targeting. Hence, AdVeil can be deployed (with off-the-shelf cloud infrastructure) and *still* result in net profit gains for the Broker.

8 RELATED WORK

There exists a large body of works on privacy-preserving advertising [9, 14, 43, 45–47, 52, 73, 90, 91, 96, 97]. In this section we focus on providing a detailed comparison to other systems that cover the complete advertising pipeline and leave discussion of the less similar works to Appendix A.

Jules [52] is an early, theoretical proposal for supporting privacy-preserving targeted advertising on the internet. It is primarily of historical interest, but shares some similarities to AdVeil. Specifically with respect to the use of PIR (which is applied to *delivery* of ads). We show in Section 7, that PIR is not practical for ad delivery.

[17]Note that Figure 4a includes network latency in addition to processing time.

Privad [45] provides a targeting model based on broad interest categories that are narrowed locally by the client. Privad relies on a *centralized* anonymizing proxy, referred to as the Dealer, to provide user privacy and enforce fraud prevention. The Dealer is assumed not to collude with the Broker and provides user privacy by mediating *all user communication*. Thus, the Broker learns the contents, but not the origin, of each request. To do this, the Dealer must sustain the load of the entire ad network as a non-colluding third party. Privad cannot trivally solve this issue by replacing the Dealer with an existing, distributed anonymizing proxy as this would leave their system without fraud prevention.

OblivAd [9] relies heavily on a Trusted Execution Environment (TEE) to provide user privacy. The TEE is single-handedly responsible for *all stages of the advertising pipeline*, from ad targeting to unlinkability of reports and fraud prevention. To prevent the Broker from learning which ads are delivered to which clients, the TEE uses ORAM [85] to retrieve ads from the ad database. For reporting, users again encrypt their responses to the TEE which batches and shuffles them to hide timing information.

ObliviAd provides strong privacy guarantees only because of its reliance on the TEE and ORAM, both of which have major practical drawbacks. ORAM can only serve a single client request at a time; parallel requests require an equivalent number of separate instances of the ORAM scheme causing linear storage blowup in the number of concurrent client requests [8, 85]. TEEs have seen a series of powerful attacks since ObliviAd was published [16, 21, 28, 93]. As a result, we do not believe that either issue is surmountable with today's instantiations of TEEs or ORAM [8, 28, 85].

Adnostic [90] primarily focuses on the targeting and reporting stages. Similarly to Privad, Adnostic performs targeting locally on the client. However, Agnostic only uses *contextual* features during targeting and does not make an effort to hide which ads are delivered to a user. Adnostic attempts to prevent the Broker from linking the delivery and reporting phases by aggregating reports using homomorphic encryption and zero-knowledge proofs rather than revealing them individually. Decryption of aggregate reports is handled by a TTP that checks the size of reports prior to decryption to ensure that only large groups of users are reported on. However, Adnostic makes no effort to hide users' browsing or click behavior from the Broker, leading us to agree with Privad's description of their Broker model as "honest-and-*not*-curious" [45].

themis [73] is the Brave browser's contribution to private targeted advertising. It attempts to replace the role of the Broker with a permissioned blockchain, run by Publishers or foundations such as the EFF. themis additionally supports payment to users for their interaction with ads. Privacy for payments, metrics, and auditing of the blockchain is based on a Proof of Authority protocol [1].

The attempted removal of the Broker, auditability, and user compensation are all interesting directions for the future of private advertising. Unfortunately, themis achieves them by entirely offloading both targeting and delivery to clients. Users must download both the targeting model and the *entire* database of ads and ad features to their local device. Fetching fewer ads is not trivial for themis due to the timing information revealed on its blockchain. We do not believe this to be insurmountable, but it would require careful analysis of

Table 9: Comparison of AdVeil to related work on targeted advertising.

	Privacy		Correctness			Scalability		
	Trusted Third Party (TTP)	Targeting Data Sent to Broker	Targeting Accuracy	Fraud Prevention	Report Granularity	TTP Workload	TTP Availability	Simultaneous Requests
Adnostic [90]	Decryption Oracle	Contextual	Contextual+	Limited [i]	Aggregate	≪ Broker	≪ Broker	✓
ObliviAd [9]	TEE	None[ii]	Full Targeted	Full	Individual	N/A [iii]	N/A	✗
Privad [45]	Dealer	Broad Categories	Limited Targeted	Full	Individual	≈ Broker	= Broker [iv]	✓
THEMIS [73]	PoA Blockchain	None	Full Targeted	Full	Individual	≈ Broker [v]	≈ Broker [v]	✓
TURTLEDOVE [97]	Traditional TTP	None [vi]	Limited Targeted+	N/A	Individual+	< Broker	≈ Broker [iv]	✓
AdVeil	None	None	Full Targeted	Full	Individual	N/A	N/A	✓

[i] Adnostic does not perform fraud detection beyond guaranteeing that an individual report is for a single ad.
[ii] Complete targeting data is sent to the TEE held by the Broker.
[iii] Oblivad's TTP is a TEE held by the Broker. It does not correspond to a physically or administratively separate entity.
[iv] The TTP participates in processing every targeting request, on demand. In Privad it also participates in reporting.
[v] The blockchain authorities are required to validate every report on demand and selected users are required to participate in MPC to compute metrics.
[vi] User data *is* sent to the TTP.

its security implications which we have not yet seen addressed in Brave's continued development efforts [18].

TURTLEDOVE [97] takes a similar approach to Privad where users subscribe to interest groups and can participate in local auctions to receive an ad based on these interest groups. However, unlike Privad, TURTLEDOVE does not make any assumptions about the specificity of these groups or what metadata may be included in the auction. It only requires that group size meets an unspecified k-anonymity threshold. Additionally, TURTLEDOVE requires a trusted server to support the inclusion of real time external data into the local auction. User requests to this server are *not* required to conform to any k-anonymity bound. Between the specificity of interest groups, the blind trust in this server, and the lax restrictions on reporting, it is not clear that TURTLEDOVE provides its users with any more privacy than traditional advertising.

9 DISCUSSION & CONCLUSIONS

On today's internet, there is a fundamental conflict between privacy and advertising. What we show with AdVeil is that, while the pervasive user tracking performed by ad networks is incompatible with privacy, targeted advertisements are not. AdVeil provides full compliance with technological and legislative best practices for user privacy *without* limiting the relevance of advertisements shown on the internet. Ultimately, we believe that AdVeil is a viable alternative to existing, non-private targeted advertising schemes that meets the needs of both users and ad network brokers.

Targeted Advertising can have serious negative side effects, from web tracking to discriminatory or manipulative practices [37]. In light of this, it seems challenging to advocate for even privacy-preserving targeted advertising. However, current legislative and technological efforts to prevent targeted ads or their ill effects have been markedly unsuccessful. Ad tech companies are more willing to fight lawsuits than they are to stop collecting user data [12, 54, 54, 86] and, despite anti-tracking measures, most peoples' browsers remain uniquely identifiable [20, 38, 66, 75, 102].

In AdVeil users have transparency into the ad targeting process and can avoid targeting on any features they deem sensitive. It does not require ad-tech companies to track users for targeting data, instead giving control over the collection and use of this data back to the users themselves. While AdVeil is not a complete solution to discriminatory or manipulative advertising, we believe it provides a platform for future work in this direction.

Usability motivated our decision to separate AdVeil's targeting and delivery stages. While AdVeil can deliver ads at approximately the same speed as the rest of a page's content, it cannot target them at the same rate. If targeting was performed on-demand, as delivery is, it could result in users navigating away from pages before ads finish loading. Instead, AdVeil targets ads in batches at fixed intervals such that they can be rapidly delivered on demand.

Similarly, the ad network requires that any computational overhead imposed by AdVeil be *sustainable*. While targeting and metrics are more computationally expensive than their non-private alternatives, we showed in Section 7.2.8 that AdVeil remains profitable. Additionally, as was discussed in Section 5.3, both protocols are versatile and allow the ad network its choice of targeting and measurement strategies to align with those used in practice.

Conclusions. We presented AdVeil, a system for privacy-preserving targeted advertising that addresses usability and privacy concerns of existing work. In doing so, we introduce a novel method of targeting ads that provides strong privacy guarantees for users, while ensuring targeting accuracy that is on-par with existing systems. We provide an open-source prototype implementation with an end-to-end evaluation. Our results show the practical and economic feasibility of deploying AdVeil for meeting user privacy demands and regulatory compliance, without compromising on the needs of the ad network.

10 ACKNOWLEDGEMENTS

We thank Peter Deutsch, Jules Drean, Ben Murphy, and Michael Specter for helpful discussion and feedback on this paper.

REFERENCES

[1] Consensys Quorum. https://consensys.net/quorum/.
[2] General Data Protection Regulation. https://gdpr-info.eu/, 2016.

[3] California Consumer Privacy Act of 2018. https://leginfo.legislature.ca.gov/faces/codes_displayText.xhtml?division=3.&part=4.&lawCode=CIV&title=1.81.5, 2018.

[4] Agrawal, S., and Raghuraman, S. KVaC: Key-value commitments for blockchains and beyond. In *International Conference on the Theory and Application of Cryptology and Information Security* (2020), Springer, pp. 839–869.

[5] Ali, M., Sapiezynski, P., Korolova, A., Mislove, A., and Rieke, A. Ad delivery algorithms: The hidden arbiters of political messaging. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining* (New York, NY, USA, 2021), WSDM '21, Association for Computing Machinery, p. 13–21.

[6] Andoni, A., Indyk, P., and Razenshteyn, I. Approximate nearest neighbor search in high dimensions. *arXiv preprint arXiv:1806.09823 7* (2018).

[7] Angel, S., Chen, H., Laine, K., and Setty, S. PIR with compressed queries and amortized query processing. In *2018 IEEE Symposium on Security and Privacy (SP)* (2018), IEEE, pp. 962–979.

[8] Asharov, G., Komargodski, I., Lin, W.-K., Nayak, K., Peserico, E., and Shi, E. Optorama: Optimal oblivious RAM. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques* (2020), Springer, pp. 403–432.

[9] Backes, M., Kate, A., Maffei, M., and Pecina, K. Obliviad: Provably secure and practical online behavioral advertising. In *2012 IEEE Symposium on Security and Privacy* (2012), IEEE, pp. 257–271.

[10] Backes, M., Kate, A., Manoharan, P., Meiser, S., and Mohammadi, E. AnoA: A framework for analyzing anonymous communication protocols. In *2013 IEEE 26th Computer Security Foundations Symposium* (2013), IEEE, pp. 163–178.

[11] Benaloh, J. D. C. *Verifiable secret-ballot elections*. PhD thesis, Yale University, 1987.

[12] Bodoni, S. Facebook to fight Belgian ban on tracking users (and even non-users). https://www.bloomberg.com/news/articles/2019-03-27/facebook-attack-of-belgian-order-on-user-tracking-gets-hearing, 2019. Bloomberg.

[13] Boneh, D., Bünz, B., and Fisch, B. Batching techniques for accumulators with applications to IOPs and stateless blockchains. In *Annual International Cryptology Conference* (2019), Springer, pp. 561–586.

[14] Boshrooyeh, S. T., Küpçü, A., and Özkasap, Ö. PPAD: Privacy preserving group-based advertising in online social networks. In *2018 IFIP Networking Conference (IFIP Networking) and Workshops* (2018), IEEE, pp. 1–9.

[15] Boyle, E., Gilboa, N., and Ishai, Y. Function secret sharing. In *Annual international conference on the theory and applications of cryptographic techniques* (2015), Springer, pp. 337–367.

[16] Brasser, F., Müller, U., Dmitrienko, A., Kostiainen, K., Capkun, S., and Sadeghi, A.-R. Software grand exposure: SGX cache attacks are practical. In *11th USENIX Workshop on Offensive Technologies (WOOT 17)* (Vancouver, BC, Aug. 2017), USENIX Association.

[17] Brave. Brave: The browser reimagined. https://brave.com/, 2020.

[18] Brave. Brave/bat ads and THEMIS request for comments and code (RFC&C) event. https://brave.com/themis-rfcc/, 2021.

[19] Broder, A. Z. On the resemblance and containment of documents. In *Proceedings. Compression and Complexity of SEQUENCES 1997 (Cat. No. 97TB100171)* (1997), IEEE, pp. 21–29.

[20] Budington, B. Panopticlick: Fingerprinting your web presence. USENIX Association.

[21] Bulck, J. V., Minkin, M., Weisse, O., Genkin, D., Kasikci, B., Piessens, F., Silberstein, M., Wenisch, T. F., Yarom, Y., and Strackx, R. Foreshadow: Extracting the keys to the intel SGX kingdom with transient out-of-order execution. In *27th USENIX Security Symposium (USENIX Security 18)* (Baltimore, MD, Aug. 2018), USENIX Association, p. 991–1008.

[22] Cabañas, J. G., Cuevas, Á., and Cuevas, R. Unveiling and quantifying Facebook exploitation of sensitive personal data for advertising purposes. In *27th USENIX Security Symposium (USENIX Security 18)* (Baltimore, MD, Aug. 2018), USENIX Association, pp. 479–495.

[23] Catalano, D., and Fiore, D. Vector commitments and their applications. In *International Workshop on Public Key Cryptography* (2013), Springer, pp. 55–72.

[24] Charikar, M. S. Similarity estimation techniques from rounding algorithms. In *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing* (2002), pp. 380–388.

[25] Chor, B., Gilboa, N., and Naor, M. *Private information retrieval by keywords*. Citeseer, 1997.

[26] Chor, B., Goldreich, O., Kushilevitz, E., and Sudan, M. Private Information Retrieval. In *Proceedings of IEEE 36th Annual Foundations of Computer Science* (1995), IEEE, pp. 41–50.

[27] Choudhuri, A. R., Green, M., Jain, A., Kaptchuk, G., and Miers, I. Fairness in an unfair world: Fair multiparty computation from public bulletin boards. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security* (2017), pp. 719–728.

[28] Costan, V., and Devadas, S. Intel SGX explained. *IACR Cryptol. ePrint Arch. 2016*, 86 (2016), 1–118.

[29] Cyphers, B. Google's FLoC is a terrible idea. https://www.eff.org/deeplinks/2021/03/googles-floc-terrible-idea, 2021.

[30] Danezis, G., Diaz, C., and Troncoso, C. Two-sided statistical disclosure attack. In *International Workshop on Privacy Enhancing Technologies* (2007), Springer, pp. 30–44.

[31] Danezis, G., and Serjantov, A. Statistical disclosure or intersection attacks on anonymity systems. In *International Workshop on Information Hiding* (2004), Springer, pp. 293–308.

[32] Datar, M., Immorlica, N., Indyk, P., and Mirrokni, V. S. Locality-sensitive hashing scheme based on p-stable distributions. In *Proceedings of the twentieth annual symposium on Computational geometry* (2004), pp. 253–262.

[33] Datta, A., Datta, A., Makagon, J., Mulligan, D. K., and Tschantz, M. C. Discrimination in online advertising: A multidisciplinary inquiry. *Conference on Fairness, Accountability, and Transparency 81* (2018), 20–34.

[34] Davidson, A., Goldberg, I., Sullivan, N., Tankersley, G., and Valsorda, F. Privacy pass: Bypassing internet challenges anonymously. *Proceedings on Privacy Enhancing Technologies 2018*, 3 (2018), 164–180.

[35] Dingledine, R., Mathewson, N., and Syverson, P. Tor: The second-generation onion router. Tech. rep., Naval Research Lab Washington DC, 2004.

[36] Dwork, C. Differential privacy: A survey of results. In *International conference on theory and applications of models of computation* (2008), Springer, pp. 1–19.

[37] Edelman, G. Why don't we just ban targeted advertising? https://www.wired.com/story/why-dont-we-just-ban-targeted-advertising/, 2020.

[38] Englehardt, S., and Narayanan, A. Online tracking: A 1-million-site measurement and analysis. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (New York, NY, USA, 2016), CCS '16, Association for Computing Machinery, p. 1388–1401.

[39] Garay, J., Katz, J., Maurer, U., Tackmann, B., and Zikas, V. Rational protocol design: Cryptography against incentive-driven adversaries. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science* (2013), IEEE, pp. 648–657.

[40] Gionis, A., Indyk, P., Motwani, R., et al. Similarity search in high dimensions via hashing. In *Vldb* (1999), vol. 99, pp. 518–529.

[41] Goodwin, B., and Skelton, S. K. Facebook's privacy game – how Zuckerberg backtracked on promises to protect personal data. https://www.computerweekly.com/feature/Facebooks-privacy-U-turn-how-Zuckerberg-backtracked-on-promises-to-protect-personal-data, 2019.

[42] Gorbunov, S., Reyzin, L., Wee, H., and Zhang, Z. Pointproofs: aggregating proofs for multiple vector commitments. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security* (2020), pp. 2007–2023.

[43] Green, M., Ladd, W., and Miers, I. A protocol for privately reporting ad impressions at scale. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (2016), pp. 1591–1601.

[44] Groce, A., Katz, J., Thiruvengadam, A., and Zikas, V. Byzantine agreement with a rational adversary. In *International Colloquium on Automata, Languages, and Programming* (2012), Springer, pp. 561–572.

[45] Guha, S., Cheng, B., and Francis, P. Privad: Practical privacy in online advertising. In *USENIX conference on Networked systems design and implementation* (2011), pp. 169–182.

[46] Helsloot, L. J., Tillem, G., and Erkin, Z. AHEad: privacy-preserving online behavioural advertising using homomorphic encryption. In *2017 IEEE Workshop on Information Forensics and Security (WIFS)* (2017), IEEE, pp. 1–6.

[47] Helsloot, L. J., Tillem, G., and Erkin, Z. BAdASS: Preserving privacy in behavioural advertising with applied secret sharing. In *International Conference on Provable Security* (2018), Springer, pp. 397–405.

[48] Iab. Iab internet advertising revenue report. *Interactive Advertising Bureau. New York* (2020).

[49] Imana, B., Korolova, A., and Heidemann, J. Auditing for discrimination in algorithms delivering job ads. WWW '21, Association for Computing Machinery, p. 3767–3778.

[50] Indyk, P., and Motwani, R. Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing* (1998), pp. 604–613.

[51] Joseph, S. 'contextual targeting is going to be the new black': As IDFA restrictions loom, advertisers brace for the fallout. https://digiday.com/media/contextual-targeting-is-going-to-be-the-new-black-as-idfa-restrictions-loom-advertisers-brace-for-the-fallout/, 2020.

[52] Juels, A. Targeted advertising... and privacy too. In *Cryptographers' Track at the RSA Conference* (2001), Springer, pp. 408–424.

[53] Juvekar, C., Vaikuntanathan, V., and Chandrakasan, A. GAZELLE: A low latency framework for secure neural network inference. In *27th USENIX Security Symposium (USENIX) Security 18)* (2018), pp. 1651–1669.

[54] Kaplan, M. Facebook and Google are already facing lawsuits under new data rules. https://money.cnn.com/2018/05/25/technology/gdpr-compliance-facebook-google/index.html, 2018.

[55] Kaspar, K., Weber, S. L., and Wilbers, A.-K. Personally relevant online advertisements: Effects of demographic targeting on visual attention and brand evaluation. *PloS one 14*, 2 (2019), e0212419.

[56] Kreuter, B., Lepoint, T., Orrù, M., and Raykova, M. Anonymous tokens with private metadata bit. In *Annual International Cryptology Conference* (2020), Springer, pp. 308–336.

[57] Kumar, N., Rathee, M., Chandran, N., Gupta, D., Rastogi, A., and Sharma, R. Cryptflow: Secure tensorflow inference. In *2020 IEEE Symposium on Security and Privacy (SP)* (2020), IEEE, pp. 336–353.

[58] Kushilevitz, E., and Ostrovsky, R. Replication is not needed: Single database, computationally-private information retrieval. In *Proceedings 38th annual symposium on foundations of computer science* (1997), IEEE, pp. 364–373.

[59] Laurie, B. Certificate transparency. *Communications of the ACM 57*, 10 (2014), 40–46.

[60] Libert, B., and Yung, M. Concise mercurial vector commitments and independent zero-knowledge sets with short proofs. In *Theory of Cryptography Conference* (2010), Springer, pp. 499–517.

[61] LLP, R. . C. California's Consumer Privacy Rights Act: Opt-out rights and data profiling. https://www.natlawreview.com/article/california-s-consumer-privacy-rights-act-opt-out-rights-and-data-profiling, 2021.

[62] Loesing, K., Murdoch, S. J., and Dingledine, R. 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)* (January 2010), LNCS, Springer.

[63] Mallesh, N., and Wright, M. The reverse statistical disclosure attack. In *International Workshop on Information Hiding* (2010), Springer, pp. 221–234.

[64] Match2One. Top banner sizes: The most effective banners of 2020. https://www.match2one.com/blog/standard-banner-sizes/, July 2020.

[65] Mathewson, N., and Dingledine, R. Practical traffic analysis: Extending and resisting statistical disclosure. In *International Workshop on Privacy Enhancing Technologies* (2004), Springer, pp. 17–34.

[66] Mathur, A., Vitak, J., Narayanan, A., and Chetty, M. Characterizing the use of browser-based blocking extensions to prevent online tracking. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)* (Baltimore, MD, Aug. 2018), USENIX Association, pp. 103–116.

[67] Matz, S. C., Kosinski, M., Nave, G., and Stillwell, D. J. Psychological targeting as an effective approach to digital mass persuasion. *Proceedings of the National Academy of Sciences 114*, 48 (2017), 12714–12719.

[68] Media, A. Digital ad specs. https://www.advancemediany.com/wp-content/uploads/2018/11/DigitalAdSpecs_20181108.pdf, May 2021.

[69] Melara, M. S., Blankstein, A., Bonneau, J., Felten, E. W., and Freedman, M. J. CONIKS: Bringing key transparency to end users. In *24th USENIX Security Symposium (USENIX Security 15)* (2015), pp. 383–398.

[70] Merkle, R. C. A digital signature based on a conventional encryption function. In *Conference on the theory and application of cryptographic techniques* (1987), Springer, pp. 369–378.

[71] Nakamoto, S. Bitcoin: A peer-to-peer electronic cash system. Tech. rep., Manubot, 2019.

[72] O'Flaherty, K. This is why people no longer trust Google and Facebook with their data. https://www.forbes.com/sites/kateoflahertyuk/2018/10/10/this-is-why-people-no-longer-trust-google-and-facebook-with-their-data, 2018.

[73] Pestana, G., Querejeta-Azurmendi, I., Papadopoulos, P., and Livshits, B. THEMIS: Decentralized and trustless ad platform with reporting integrity, 2020.

[74] Plane, A. C., Redmiles, E. M., Mazurek, M. L., and Tschantz, M. C. Exploring user perceptions of discrimination in online targeted advertising. In *26th USENIX Security Symposium (USENIX Security 17)* (Vancouver, BC, Aug. 2017), USENIX Association, pp. 935–951.

[75] Pugliese, G., Riess, C., Gassmann, F., and Benenson, Z. Long-term observation on browser fingerprinting: Users' trackability and perspective. *Proceedings on Privacy Enhancing Technologies 2020*, 2 (01 Apr. 2020), 558 – 577.

[76] Rajaraman, A., and Ullman, J. D. *Mining of massive datasets*. Cambridge University Press, 2011.

[77] Rathee, D., Rathee, M., Kumar, N., Chandran, N., Gupta, D., Rastogi, A., and Sharma, R. Cryptflow2: Practical 2-party secure inference. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security* (2020), pp. 325–342.

[78] Schraeder, P. Contextual advertising leads the way in a post cookie reality. https://adage.com/article/gumgum/contextual-advertising-leads-way-post-cookie-reality/2293976, 2020.

[79] Silde, T., and Strand, M. Anonymous tokens with public metadata and applications to private contact tracing. *IACR Cryptol. ePrint Arch. 2021* (2021), 203.

[80] Singer, N. What you don't know about how Facebook uses your data. https://www.nytimes.com/2018/04/11/technology/facebook-privacy-hearings.html, 2018.

[81] Slaney, M., Lifshits, Y., and He, J. Optimal parameters for locality-sensitive hashing. *Proceedings of the IEEE 100*, 9 (2012), 2604–2623.

[82] SPEEDTEST, B. O. Speedtest global index. Tech. rep., citirano 2020-1-30, dostupno na: https://www. speedtest. net/global-index ..., 2019.

[83] Springborn, K., and Barford, P. Impression fraud in on-line advertising via pay-per-view networks. In *22nd USENIX Security Symposium (USENIX Security 13)* (2013), pp. 211–226.

[84] Statista. Social media advertising cost-per-mille (CPM) worldwide from 2nd quarter 2018 to 4th quarter 2019. https://www.statista.com/statistics/873631/social-media-advertising-cpm/, 2020.

[85] Stefanov, E., van Dijk, M., Shi, E., Fletcher, C., Ren, L., Yu, X., and Devadas, S. Path ORAM: An extremely simple oblivious RAM protocol. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security* (New York, NY, USA, 2013), CCS '13, Association for Computing Machinery, p. 299–310.

[86] Stempel, J. Google faces $5 billion lawsuit in U.S. for tracking "private" internet use. https://www.reuters.com/article/us-alphabet-google-privacy-lawsuit/google-faces-5-billion-lawsuit-in-u-s-for-tracking-private-internet-use-idUSKBN23933H, 2020.

[87] Sweeney, L. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems 10*, 05 (2002), 557–570.

[88] Team, T., and Speculations, G. Is Google advertising revenue 70%, 80%, or 90% of Alphabet's total revenue? https://www.forbes.com/sites/greatspeculations/2019/12/24/is-google-advertising-revenue-70-80-or-90-of-alphabets-total-revenue/, 2019-12-24.

[89] Tomescu, A., Xia, Y., and Newman, Z. Authenticated dictionaries with cross-incremental proof (dis) aggregation. *IACR Cryptol. ePrint Arch. 2020* (2020), 1239.

[90] Toubiana, V., Narayanan, A., Boneh, D., Nissenbaum, H., and Barocas, S. Adnostic: Privacy preserving targeted advertising. In *Proceedings Network and Distributed System Symposium* (2010).

[91] Tran, M.-D., Acs, G., and Castelluccia, C. Retargeting without tracking. *arXiv preprint arXiv:1404.4533* (2014).

[92] Tulabandhula, T., Vaya, S., and Dhar, A. Privacy-preserving targeted advertising. *arXiv preprint arXiv:1710.03275* (2017).

[93] van Schaik, S., Kwong, A., Genkin, D., and Yarom, Y. SGAxe: How SGX fails in practice. https://sgaxe.com/files/SGAxe.pdf, 2020.

[94] Varvello, M., Azurmendi, I. Q., Nappa, A., Papadopoulos, P., Pestana, G., and Livshits, B. VPN0: A privacy-preserving decentralized virtual private network. *arXiv preprint arXiv:1910.00159* (2019).

[95] Wang, F., Yun, C., Goldwasser, S., Vaikuntanathan, V., and Zaharia, M. Splinter: Practical private queries on public data. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)* (2017), pp. 299–313.

[96] Web Incubator CG. FLOC. https://github.com/WICG/floc, 2021.

[97] Web Incubator CG. TURTLEDOVE. https://github.com/WICG/turtledove, 2021.

[98] Wei, M., Stamos, M., Veys, S., Reitinger, N., Goodman, J., Herman, M., Filipczuk, D., Weinshel, B., Mazurek, M. L., and Ur, B. What Twitter knows: Characterizing ad targeting practices, user perceptions, and ad explanations through users' own twitter data. In *29th USENIX Security Symposium (USENIX Security 20)* (Aug. 2020), USENIX Association, pp. 145–162.

[99] Wilander, J. Privacy preserving ad click attribution for the web. https://webkit.org/blog/8943/privacy-preserving-ad-click-attribution-for-the-web/, 2019.

[100] Wolinsky, D. I., Syta, E., and Ford, B. Hang with your buddies to resist intersection attacks. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security* (2013), pp. 1153–1166.

[101] Wright, C., and Varia, M. Crypto crumple zones: Enabling limited access without mass surveillance. In *2018 IEEE European Symposium on Security and Privacy (EuroS&P)* (2018), IEEE, pp. 288–306.

[102] Yen, T.-F., Xie, Y., Yu, F., Yu, R. P., and Abadi, M. Host fingerprinting and tracking on the web: Privacy and security implications. In *NDSS* (2012), vol. 62, p. 66.

[103] Zantout, B., Haraty, R., et al. I2P data communication system. In *Proceedings of ICN* (2011), Citeseer, pp. 401–409.

[104] Zhang, R., and Cui, Y. Similarity function in online advertising bid optimization, Aug. 4 2011. US Patent App. 12/698,463.

[105] Zhu, X., Tao, H., Wu, Z., Cao, J., Kalish, K., and Kayne, J. *Fraud prevention in online digital advertising*. Springer, 2017.

A EXTENDED RELATED WORK

FLoC [96] focuses specifically on the targeting phase of the advertising pipeline. It assigns users to groups based on their browsing behavior – users with similar browsing habits and, presumably, interests are assigned to the same group. FLoC users receive ads based on a group identifier rather than a personal feature vector. The proposal believes this to be secure because the number of users per-group is large and the total number of groups is small, indicating that no one group can be too specific to an individual user. Unfortunately, FLoC considers "thousands" of users to be a sufficiently large group size and the group name "43A7" to be a short group identifier. Of course, "thousands" is a *very* small group size compared to the 4.66 billion active internet users and 1,679,616 different groups can be supported, assuming 4 digit alphanumeric representation for group identifiers. Thus, it is unclear what, if anything, sending this group name instead of a feature vector hides about users of FLoC as groups could be highly specific to users' private features.

PPAD [14] is a group-based ad targeting system where the privacy for users is guaranteed relative to the *group* that a user belongs to. Each user is assigned to a group based on their attributes and hence reveals *coarse grained* interests of users. The Broker and a semi-trusted third party run a variant of a private-set intersection protocol (using shares of user-provided feature vectors) to determine which ads should be displayed to a group. The advantage of PPAD is that the Broker can target and deliver ads at a group level and can evaluate set intersection while the user is "offline" in order to serve ads when the user goes back online. However, the model of PPAD provides a tradeoff between targeting accuracy and privacy. Specifically, the more fine-grained the group selection, the more privacy leakage it incurs; PPAD does not provide an analysis of this leakage.

BAdASS [47] (and its precursor AHEAd [46]) are designed for online targeting and do not delve into other aspects of the advertising ecosystem such as reporting and fraud prevention. BAdASS uses a multi-party computation protocol executed between a group of honest-but-curious parties. Moreover, BAdASS requires splitting trust among a set of Demand-Side Platforms (DSPs) which manage content targeting in an ad network. Even a single malicious DSP can disrupt the correctness of the protocol (or worse yet de-anonymize the user) unless expensive zero-knowledge proofs are added to prevent deviations from protocol.

Tran *et al.* [91] is designed to address the *retargeting* aspect of online advertising. Their primary assumption is that the Broker (or ad exchange) will not collude with retargeting services. Clients engage in a protocol between the Broker and retargeter to fetch ads for products they have previously shown interest in (e.g., by adding a product to a shopping cart). The retargeter learns which ad was displayed but not which user requested it, while the Broker learns which user requested the ad but not which ad was retrieved via the retargeter. User privacy is ensured if the Broker and retargeter do not collude.

Tulabandhula *et al.* [92] propose a collection of functions for privacy-preserving association rule mining and recommendations. Their protocols work between a client and server, where the client stores the feature vector locally. Their results can be applied to AdVeil in the targeting process and may be useful for certain deployments.

AdScale [43] improves the reporting scheme proposed in Adnostic, but does not address other aspects of the pipeline. Users in Adscale respond with homomorphically encrypted reports that are aggregated by the Broker, but can only be decrypted by a designated trusted third party (TTP). This is intended to ensure that neither the Broker nor the TTP are individually capable of learning the plaintext responses of a single user. As a consequence, only aggregate information can be used for billing and targeting purposes.

B PRIVATE KEYWORD ADVERTISING

While we present a protocol for privacy-preserving ad selection using nearest neighbor search (e.g., for selecting ads based on user interests), a different category of advertising involves displaying ads to users based on *keywords*. For example, searching for "hotels in madagascar" in DuckDuckGo displays an ad by `booking.com` for hotel deals in Madagascar as well as an ad for cheap plane tickets. These ads are served by matching keywords in the search query ("hotel" and "madagascar") to a "bag of words" associated with each ad. Our protocol for privacy preserving nearest neighbor search can easily be adapted to match a set of keywords.