Journal of Communications and Information Networks
Vol.1, No.4, Dec. 2016
DOI: 10.11959/j.issn.2096-1081.2016.054

Review paper

Towards privacy-preserving user targeting

JIANG Jinghua[1,2†], ZHENG Yifeng[2†], SHI Zhenkui[2], YAO Jing[1,2], WANG Cong[2,3], GUI Xiaolin[1]

1. Xi'an Jiaotong University, Xi'an 710049, China
2. City University of Hong Kong, Hong Kong 999077, China
3. City University of Hong Kong Shenzhen Research Institute, Shenzhen 518057, China

Abstract: User targeting via behavioral analysis is becoming increasingly prevalent in online messaging services. By taking into account users' behavior information such as geographic locations, purchase behaviors, and search histories, vendors can deliver messages to users who are more likely to have a strong preference. For example, advertisers can rely on some ad-network for distributing ads to targeted users. However, collecting such personal information for accurate targeting raises severe privacy concerns. In order to incentivize users to participate in such behavioral targeting systems, addressing the privacy concerns becomes of paramount importance. We provide a survey of privacy-preserving user targeting. We present the architectures of user targeting, the security threats faced by user targeting, and existing approaches to privacy-preserving user targeting. Some future research directions are also identified.

Key words: privacy, ad-network, user targeting, advertising, coupon delivery

Citation: JIANG J H, ZHENG Y F, SHI Z K, et al. Towards privacy-preserving user targeting[J]. Journal of communications and information networks, 2016, 1(4): 22-32.

1 Introduction

User targeting via behavioral analysis, also known as behavioral targeting, is becoming increasingly prevalent in online messaging services[1]. It brings benefits to both vendors and users. From the vendor's perspective, user targeting is advantageous as it allows them to deliver messages to users who are more likely to have a strong preference. From the user's perspectives, user targeting allows them to enjoy more personalized services and get less exposure to messages that are not pertaining to their interests.

Roughly speaking, there are two types of typical applications in the current practice of user targeting. The first one is known as targeted advertising, where advertisers want their ads to be delivered to users who are likely to have strong interests, via taking into account their behavior information. A realistic example of targeted advertising systems is Google AdWords, which delivers personalized ads

Manuscript received Sept. 26, 2016; accepted Nov. 11, 2016
This work is supported by the National Natural Science Foundation of China (Nos. 61572412, 61472316), the Research Grants Council of Hong Kong (Nos. CityU 138513, CityU 11276816), the Innovation and Technology Commission of Hong Kong under ITF Project (No. ITS/307/15), an AWS Education Research Grant, and the Science and Technology Project of Shaanxi Province (Nos. 2016ZDJC-05, 2013SZS-16).

† The first two authors contributed equally to this paper.

based on users' search items. The other kind of user targeting application is targeted coupon delivery[2,3], in which vendors aim to deliver targeted coupons to certain users who are likely to become loyal routine customers.

Despite the well-understood benefits, user targeting also raises severe privacy concerns. In order to target accurately, vendors need to collect adequate behavior information about users, such as their geographic locations, purchase behaviors, and search histories. Exposing such personal information may easily violate user privacy[2,4]. Therefore, addressing the privacy concerns of users becomes crucial to incentivize them to participate in user targeting systems.

When it comes to specific applications, additional security challenges may be further posed. In particular, targeted coupon delivery raises the following additional security requirements. First, coupons are required to be delivered only to the eligible users so as to prevent coupon exploits. The strategy of simply pushing down all the coupons in clear to the user's device and asking it to select the ones that the user is eligible for does not work well. This simple strategy exposes the coupons to malicious users who may attempt to get discounts that they are not eligible for. Second, targeted coupon delivery systems should ensure that the vendor's targeting strategy (i.e., the coupon's eligibility requirements) is well protected against non-eligible users[2,3]. Therefore, in order to defend against various potential threats to user targeting systems, security must be embedded in the system design from the very beginning.

In this paper, we survey the literature related to the typical types of applications of privacy-preserving user targeting, i.e., private targeted advertising and private targeted coupon delivery. We present the architectures of user targeting and the security threats in different applications, and survey existing solutions that enable user targeting with privacy assurance. This paper is organized as follows. Section 2 presents the architecture and security threats of targeted advertising, and describes existing work on privacy-preserving targeted advertising. Section 3 presents the architecture and security threats of targeted coupon delivery, and describes existing work on privacy-preserving targeted coupon delivery. Section 4 discusses some future research directions. Section 5 concludes the whole paper.

2 Private user targeting in advertising

2.1 Problem statement

System Architecture. Fig.1 illustrates the architecture of targeted advertising. There are four primary parties: the publisher, the advertiser, the ad-network, and the user. The role of each party is described as follows. The publisher owns pages or applications and is willing to place ads in its pages or applications, and expects to be paid for this service. The advertiser expects their ads to be precisely delivered to potential purchasers (users) across different pages or applications, and is willing to pay for this service. The user wants to enjoy personalized ad delivery services so that she can receive ads pertaining to her interests. She may click on the ads displayed to her. The ad-network connects advertisers, publishers, and users. It collects ads and metadata from advertisers and places ads on registered pages or applications of publishers. Besides, the ad-network collects the view/click/action reports from users for billing advertisers and sharing part of its revenue with publishers.

The workflow between the parties in targeted advertising can be described as follows. Advertisers submit ads to the ad-network with the bid information and the attributes for user targeting, such as demographics and keywords of interests. When a user browses a website or runs an application that has some assigned screen real estate for ad display, the web-site/application would enable the ad-network to collect personal information of the user, such as demographics, location, and browsing histories. Then, the ad-network performs auctions so as to decide which ads are to delivered to the user, based on both global and user modifiers (e.g. bids, CTR (Click Through Ratio), advertiser quality, and client score)[5].

When the user's browser or application receives the targeted ads, it displays the ads to the user. Hereafter, we refer to the local browser or application run by the user as the client. And throughout the paper we will use the terms of user and client interchangeably unless otherwise stated. Meanwhile, a report of this ad view is sent to the ad-network. If the user clicks on an ad or purchases the advertised goods, a report of the corresponding action will also be sent to the ad-network. Based on the information of reports, the ad-network bills advertisers and pays publishers.

Security threats. User privacy is a critical concern that needs be carefully addressed in targeted advertising systems. Along with the workflow of targeted advertising, the types of data that a user may wish to keep private includes: local data for profiling, behavioral profile, ad view history, and ad click history. The ad-network expects to collect the private user data as much as possible for accurate targeted advertising. Also, the ad-network may collect users' data for other purposes, such as selling the information to other parties[6], so as to obtain financial incentives.

2.2 Work on private targeted advertising

As indicated by the architecture of targeted advertising, a privacy-aware targeted advertising system has two major components that involves private user data: the ad delivery component which selects the best ad for a user based on her behavioral profile; and the billing component which allows the ad-network to bill advertisers so as to get paid for ad views and clicks. Therefore, in what follows, we will first focus on describing existing work that can support these components in a privacy-preserving way, i.e., privacy-preserving ad delivery and privacy-preserving billing. Later, we will also discuss other secure components.

2.2.1 Privacy-preserving ad delivery

To enable the delivery of ads to targeted users with privacy preservation, a number of solutions have been proposed in the literature. Some representative works [4,5,7-9] are described below.

Juels'01. Juels is the first to study private targeted advertising[7]. Juels proposes a privacy-preserving targeted ad delivery scheme based on private information retrieval and mix networks. In this scheme, advertisers choose a negotiant function which assigns the best ads in the advertiser database for each type of profile. Its most sophisticated version lets each customer generate a public/private key pair and uses a bulletin board for the publishing of customers' ad requests. The advertiser server applies a mix network based on m servers, which is used to retrieve each ad with the guarantee of $(m/2, m)$-group-privacy. As the scheme relies on heavy cryptographic operations which suffers from intensive computation cost, it cannot be used to support the retrieval of ads on the fly.

Adnostic. Toubiana et al. propose an architecture named Adnostic[8], which can enable users to retrieve ads on the fly. In particular, before the user starts to visit the publisher's web page, Adnostic prefetches n ads in the browser and stores them locally. When a user browses a web page, the browser receives the list of n ads and compares it to the list of prefetched ads. If the ad chosen for display is already stored locally, the browser displays it immediately, leading

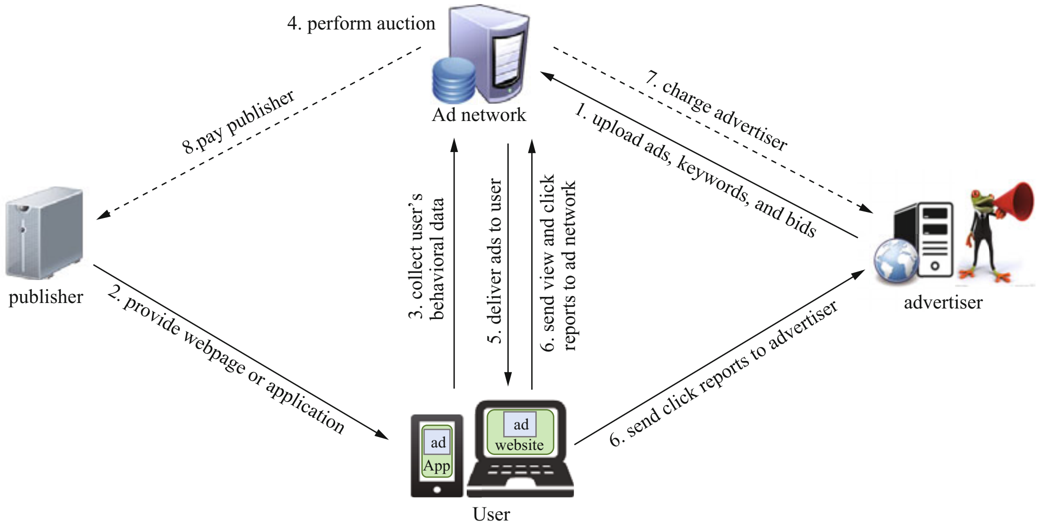
4. perform auction
Ad network
8.pay publisher
7. charge advertiser
1. upload ads, keywords, and bids
publisher
advertiser
2. provide webpage or application
3. collect user's behavioral data
5. deliver ads to user
6. send view and click reports to ad network
6. send click reports to advertiser
ad
App
ad
website
User

Figure 1 Architecture of targeted advertising

to the speed up of page display. However, regardless of whether the selected ad was prefetched or not, Adnostic always downloads the listed non-prefetched ads so as to avoid information leakage to the ad-network. Here, the choice of the parameter n is configurable. A larger value allows more precise targeting, but consumes more network bandwidth, while a smaller value results in more coarse-grained targeting, but consumes less network bandwidth. An appropriate value for n is 20, as suggested by Adnostic.

Since the browser only prefetches a small number of ads from the ad-network, a risk exists that none of the retrieved ads precisely matchs the user interest. However, to cover the spectrum of user interests, the ad-network needs to send only one ad per segment for a given interest-segmentation system. The survey conducted by Toubiana et al. found that the number of segments in existing systems is between 25 and 100, which provides an upper bound on the number of ads sent[8].

Privad. In Ref.[5], Guha, et al. present an architecture called Privad, which introduces an additional party called dealer. The dealer is mainly introduced to anonymize the client so as to prevent the ad-network from identifying the client. To prevent the dealer from accessing user's private behavioral profile, Privad encrypts the communications between the client and the dealer using the public key of the ad-network. One potential limitation of such an architecture is that it requires the dealer to always stay online, which might not be easily satisfied in practice.

ObliviAd. In Ref.[4], Backes, et al. point out that there are two drawbacks in existing anonymity networks. The first one is that they do not provide adequate performance, while ads are usually required to be displayed for users almost instantaneously. The second one is that they make the users hardly unaccountable. Therefore, they propose a new provably secure and practical online behavioral advertising architecture called ObliviAd. ObliviAd resorts to secure hardware-based PIR (Private Information Retrieval) technology, which is implemented using ORAM (Oblivious RAM) over a SC (Secure Co-processor) residing on the ad-network side. ObliviAd allows the client to retrieve ads that best match their behavioral profiles without revealing any private information to the ad-network.

In ObliviAd, to fetch an ad, a user first sends her encrypted behavioral profile to the SC. Then, SC securely selects the ads which best match the user's profile based on the algorithm specified by the ad-network. To prevent the ad-network from learning which ads are selected, they build on a state-of-the-art ORAM protocol. They also modify the adopted ORAM scheme so as to support the case of multiple entries per key-word. The selected ads are finally sent to the user in encrypted form, along with fresh electronic tokens that will be used for secure billing.

HN'12. In Ref.[9], Hardt, et al. point out that it is hardly possible to optimize the three design goals-ad relevance, privacy, and efficiency, in a single targeted advertising system. Therefore, they formalize the task of ad selection in targeted ad delivery as an optimization problem with three important variables: (1) privacy, i.e., how much information about the user's profile is shared with the ad-network, (2) communication efficiency, i.e. how few ads are sent to the user, and (3) utility, i.e., how useful the displayed ads are to the user, in terms of revenue and relevance.

In particular, their framework allows users to decide the amount of information about their sensor readings or inferred contexts that they are willing to share with the ad-network. Based on the provided information, the ad-network selects a set of ads with bounded communication overhead, and returns them to the user's client. Then, based on all the private information of the user, the client selects the most relevant ad for display. The key challenge faced by their framework is how to properly choose the set of ads sent by the ad-network and the ad displayed on the user side, in a way that can maximize utility (i.e., revenue) given constraints on efficiency (i.e., maximum communication cost) and privacy (i.e., maximum information disclosure). They show that finding the appropriate set of ads is NP hard, and propose to employ approximation techniques to solve the optimization problem.

JGSYW'15. A very recent work by Jiang et al. resorts to the PSS (Private Stream Searching) technique, which enables users to retrieve targeted ads that match their behavioral profiles in a privacy-preserving manner[10].

2.2.2 Privacy-preserving billing

As mentioned in Section 2.1, the ad-network needs to charge advertisers and share part of its revenue with publishers. Roughly speaking, there are three payment models adopted by the current practice of targeted advertising, i.e., CPC (Cost Per Click), CPM (Cost Per Impression), and CPA (Cost Per Acquisition) [11]. In the CPC and CPA models, advertisers need to pay when users click on an ad or makes a specific action (e.g., purchasing an item). In the CPM model, payment needs to be made whenever an ad is displayed to the user.

Most of existing work focuses on supporting privacy-preserving billing in the CPM/CPC model. The security goal here is to protect the ad click/view report of an individual user against the ad-network, because the ad view/click report may unexpectedly reveal user interests. On another hand, the ad-network should still be able to record the total number of views/clicks throughout a billing cycle, so as to correctly charge advertisers.

Adnostic. In Ref.[8], Toubiana, et al. propose a secure billing scheme in the CPM model for their Adnostic architecture. Adnostic adopts homomorphic encryption and zero-knowledge proof to allow the ad-network to correctly charge the corresponding advertisers, without seeing which ads are viewed by which users. It does not consider the revealing of users' ad click history as a privacy breach. The main idea of Adnostic is to let the ad-network maintain encrypted counters of ad views throughout a billing cycle, which are updated based on the homomorphic property of homomorphic encryption. At the end of a billing cycle, they resort to a trusted third party to decrypt the encrypted counters.

Privad. In Ref.[5], Guha, et al. rely on the third party dealer to anonymize the report of users' ad views and clicks, as well as other ad-initiated user activity (e.g., purchase and registration). In their scheme, a report for billing contains the ad ID, publisher ID, and type of event (view, click, etc.). Then, the report is sent to the dealer in encrypted form protected by the ad-network's public key. After anonymization, the dealer forwards the encrypted reports of users to the ad-network, which then can decrypt them for billing advertisers. Their design prevents the ad-network from linking reports to any user, and also keeps the ad-network from linking multiple reports pertaining to the same user.

ObliviAd. As mentioned before in Section 3.1, ObliviAd returns ads in encrypted form to users, along with electronic tokens that are used for billing[4]. When the presented ad is viewed or clicked by the user, the user's client sends back the token to the ad-network. The ad-network then accumulates the tokens over a predefined billing period, and sends the accumulated results to the SC. After removing duplicates (i.e., tokens with the same timestamp) and the tokens with timestamps outside of the current billing period, the SC decrypts the ads in the remaining tokens, and publishes a random permutation of the identifiers to the ad-network. At the end of a billing cycle, the ad-network distributes these identifiers to the corresponding advertisers and bills them accordingly.

KLM'12. In Ref.[12], Kodialam, et al. propose an efficient perturbation-based mechanism for protecting ad-view reports, which allows the ad-network to estimate the total number of views for each ad without learning which ads are viewed by a particular user.

In their design, the user's client first receives a set of ads from the ad-network. Then, the client selects an ad for display according to the user's profile. After this, a binary vector is generated on the user side, which indicates which ad is selected for display and which ads are not. In order to protect the information of the profile, the client then probabilistically perturbs each entry of vector. That is, it flips the value of the entry with some probability. The user then sends this perturbed vector to the ad-network.

Despite the perturbed vector sent by each user for privacy protection, the ad-network can still estimate the number of views for each ad, after collecting the perturbed vectors from all users. Based on the estimated information, the ad-network can bill advertisers without knowing individual user's profile. The key challenge here is how to accurately estimate the number of views for an ad from the perturbed vectors sent from the users. They develop an algorithm to address this challenging estimation problem, which performs one to two orders of magnitude better than standard approaches.

HTZ'15. In Ref.[13], Hua, et al. point out that, in the billing components of all existing targeted advertising systems, the ad-network exclusively determines the payment to get from advertisers and the revenue to share with publishers. Such unfairness would enable a malicious ad-network to overcharge advertisers or underpay publishers. Moreover, as bills cannot be justified by the ad-network, malicious advertisers may deny actual views/clicks to ask for refunds. Also, malicious publishers may claim clicks that actually do not exist, so as to demand extra revenue shares.

To address the potential threats from the malicious ad-network, advertisers and publishers, Hua, et al. propose to reverse the priority between these parties[13]. Specifically, when users click on ads, the client sends click reports in encrypted form to advertisers and publishers, each of which forwards the reports to the ad-network after checking, anonymizing and signing them. Here, the click reports are encrypted by the public key of the ad-network, so that they are protected against advertisers and publishers. As the billing is now based on the click reports signed by advertisers/publishers themselves, the problem of unfairness as mentioned before is addressed.

Additionally, Hua, et al. consider that (1) malicious advertiser may underreport actual ad clicks (i.e., drop click reports from clients) to reduce their bills, and (2) malicious publishers may overreport ad clicks (i.e., add fake click reports) to drive up their revenue shares. They also propose some effective detection methods to defend against these attacks, so as to ensure the safeguard the interests of the ad-network.

2.2.3 Other secure components

Private statistics gathering. In targeted advertising systems, the ad-network may require various historical global statistics on ad clicks/views. For example, to potentially maximize the revenue, the ad-network needs the CTR information in different profiles and contexts for ad selection during ad delivery. However, collecting individual user's historical statistics raises privacy concerns. In the literature, some solutions have been proposed in Refs.[9,13-20] to tackle this problem, enabling private statistics gathering. The goal of these solutions is to enable to the ad-network to learn the global statistics while preventing it from getting any individual information.

In Ref.[13], Hua, et al. propose an effective scheme for privacy-preserving statistics gathering. In their scheme, when the ad-network wants to estimate global statistics, it sends a query to a number of online clients. Then, each client sends its encrypted statistics to the advertiser, which accumulates and forwards them to the ad-network after proper anonymization. As the statistics are encrypted under the ad-network's public key, the ad-network can then decrypt them for aggregation.

Despite the effectiveness, the proposed scheme does not provide differential privacy guarantees for users[16]. That is, the ad-network might be able to reveal a specific user's response by comparing the aggregation results of different queries. To support statistics gathering with differential privacy guarantees, there is also a line of work proposing viable approaches[9, 14-20].

Click-fraud defense. In targeted advertising systems, un-scrupulous or compromised users may attempt to launch click-fraud attacks. Click-fraud attacks refer to that some users click on ads for the purpose of attacking one or more parties in the ad network[5, 13]. For example, a malicious publisher may hire some persons to click on ads displayed on its websites to drive up its revenue from the ad-network. In current targeted advertising systems, there are not perfect solutions to the defense against click-fraud attacks. The common approaches adopted in practice include: Per-User Thresholds, Blacklist, Honeyfarms, Historical Statistics, Premium Clicks, and Bait Ads[5].

On another hand, protection of user privacy makes click-fraud more challenging as clients might be hidden from the ad network. In Privad, Guha et al. point out that the dealer can assist the ad-network in detecting click-fraud[5]. In particular, the dealer measures the number of ad views/clicks from clients. Meanwhile, the ad-network inspects overall click behavior for advertisers and publishers, and informs the dealer of suspected reports. The dealer then traces the reports back to the client, and may trigger an action of blocking of that client's subsequent reports. Similarly, Hua et al. resort to the advertisers to assist the ad-network to detect click-fraud attacks[13].

3 Private user targeting in coupon delivery

3.1 Problem statement

System architecture. The architecture of targeted coupon delivery is illustrated in Fig.2. At the core, it

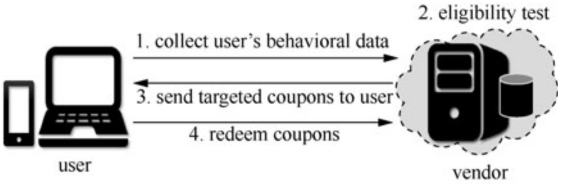
2. eligibility test
1. collect user's behavioral data
3. send targeted coupons to user
4. redeem coupons
user
vendor

Figure 2 Architecture of targeted coupon delivery

includes two primary parties: the user and the vendor. The user maintains a private behavioral profile on her mobile device, and wants to enjoy personalized coupon delivery service. The vendor wants to perform accurate user targeting via taking into account user behavior, and deliver targeted coupons only to eligible users whose behavioral profiles accurately satisfy the targeting strategies of the coupons.

Behavior Encoding. According to existing work[2, 3], both the vendor's targeting strategy and the user's behavioral profile can be represented as vectors. In particular, the behavioral profile $\boldsymbol{u}$ of a user could be represented as an n-dimensional vector, i.e., $\boldsymbol{u}=\{u_i\}_{i=1}^{n}$, where each element u_i can be either an integer or real number that denote representative statistics of different kinds of user behavior over a certain amount of time[2]. Likewise, the targeting strategy $\boldsymbol{v}$ of a coupon can also be represented as an n-dimensional vector, i.e., $\boldsymbol{v}=\{v_i\}_{i=1}^{n}$, which characterizes the user behavior that the vendor targets.

Eligibility requirement. Different encodings of the user's behavioral profile and the vendor's targeting strategy could lead to different coupon eligibility requirements. Roughly speaking, there are two cases for determining user eligibility.

- The entries in $\boldsymbol{v}$ approximately match entries in $\boldsymbol{u}$. In this case, the vendor distributes coupons by relying on a series of predictive features, which are numerical values derived from user's daily behavior stream. The similarity between $\boldsymbol{v}$ and $\boldsymbol{u}$ is measured by some kind of distance metric (e.g. Euclidean distance and cosine distance), which is dependent on the specific application scenarios.
- The entries in $\boldsymbol{v}$ exactly match entries in $\boldsymbol{u}$. In this case, the vendor distributes coupons by relying on a series of deterministic rules. For example, the vendor may only care about whether the user has been to a certain local retail store, or has been the vendor's loyalty program member to offer him a coupon. Different from the former case which requires distance calculation, this case only requires equality testing to determine whether a user is eligible for coupons.

Security threats. A secure targeted coupon delivery system should be able to address the protection from the user and vendor sides. Regarding user protection, any information of the user' personal behavioral profile should be kept confidential against the vendor throughout the targeted coupon delivery process, unless a coupon is redeemed. This means that the behavioral profile and the eligibility status of the user are both protected against the vendor during the process of targeted coupon delivery. From the vendor's perspective, its coupon and targeting strategy should be kept private against non-eligible users throughout the delivery process. More precisely, it is required that a user either learns her eligibility for a particular coupon, or learns nothing beyond her non-eligibility.

3.2 Work on private targeted coupon delivery

There are some initial efforts towards privacy-preserving targeted coupon delivery, which are described below.

PiCoDa. Partridge, et al. propose a privacy-preserving coupon delivery architecture named PiCoDa[2]. Their architecture supports both cases of eligibility testing for targeted coupon delivery in a privacy-preserving way.

In the case of approximate matching, they propose a non-interactive design, in which the user's behavioral profile never leaves the local device. In particular, they resort to the technique of LSH (Locality-Sensitive Hashing) for efficient approximate matching. LSH is a well-studied effective clustering algorithm, which hashes high-dimensional data points in such a way that the close ones collide with much higher probability than distant ones. That is, two vectors can be deemed as similar if their LSH values are equal.

In their non-interactive design, coupons in encrypted form are pushed down to the user's device. It is ensured that users can obtain the key for coupon decryption, if their behavioral profiles have the LSH value matched with that of the vendor's targeting profile. Despite the obvious merits, their design inevitably exposes the coupons to a portion of non-eligible users, due to the existence of false positives in LSH.

In the case of exact matching, they relax their stringent constraints and propose an interactive design, which requires the user to interact with the vendor during the process of coupon delivery. In particular, they resort to the cryptographic primitive known as PAKE (Password Authenticated Key Exchange), which allows two parties to agree on a session key if and only if they share a short secret. In their proposed PAKE protocol, the secret input of the user is the hash of her behavioral profile, while the secret input of the vendor is the hash of the targeting strategy. The session key derived from the PAKE protocol is then used to encrypt coupons. Consequently, if the user's behavioral profile exactly matches the targeting strategy, she can obtain the same session key from the PAKE protocol, and thus is able to decrypt the encrypted coupon.

RU'14. In Ref.[3], Rane, et al. propose a different approach to supporting eligibility testing in the case of approximate matching. Their scheme relies on the technique of fuzzy commitment, and is designed to be non-interactive. Fuzzy commitment is originally a method of supporting access control without storing a biometric in the clear at device. In fuzzy commitment, a codeword is generated from a secret key and then perturbed by the user's biometric at enrollment. Therefore, the key is bound to the biometric. A user can then extract the key and gain access, if and only if her test biometric approximately matches the enrollment biometric. Fuzzy commitment can be built from error correcting codes.

Their design applies error correcting codes to the user's behavioral profile and the vendor's targeting profile. Only if the user's behavioral profile approximately matches the targeting strategy, she can extract the key for coupon decryption. Although quite effective, their design reveals some information about the vendor's targeting profile to non-eligible users. Moreover, it requires binary entries of both the user's behavioral profile and the vendor's targeting strategy, which might not be always satisfied in practice.

4 Research directions

Mobile-oriented performance. Most of existing work focuses on online Web advertising. With the ubiquity of mobile devices, the market of mobile targeted advertising is becoming increasingly significant. As mobile devices are usually equipped with limited resources, especially battery and band-width, directly applying existing private targeted advertising techniques may not satisfy the stringent performance requirements of mobile devices. Therefore, designing targeted advertising systems with mobile-oriented performance is of practical importance[9-11, 21].

Malicious Ads. Media-rich ads on mobile devices usually contain JavaScript, image or video. These ads usually need to access the external storage, which is a shared cache where multiple apps store their files. Although the enforcement of the same origin policy prevents confined ads from accessing the external-storage files of other apps, it does not necessarily keep them from learning the existence of a file with a particular name. Consequently, malicious ads may use the existence of a file to infer sensitive information about users. For example, they may infer the user's gender preference in dating, and the user's social circle. To address the emerging threats of malicious ads, it becomes essential to re-design the mobile advertising software stack, so as to provide an isolated execution environment for the entire functionality required by media-rich ads[22, 23].

Mobile ad fraud. Ad fraud has been extensively studied in the context of online web advertising. However, it has gone largely unstudied in the context of mobile advertising. As attackers can simply distribute their malicious applications through mobile application markets, mobile devices become a lucrative target for those who launch ad fraud attacks professionally. Therefore, designing fraud detection mechanisms that can accurately and quickly identify fraud, and scale to thousands of visually complex apps is an important research direction[24, 25].

5 Conclusion

In this paper, we surveyed the literature related to private targeted advertising and private targeted coupon delivery, which are the typical applications of privacy-preserving user targeting. We started with the presentation of private targeted advertising, where we described the architecture of targeted advertising, the security threats, and existing solutions proposed for private targeted advertising. We then moved to the presentation of private targeted coupon delivery, where we described the architecture of targeted coupon delivery, behavior encoding, eligibility requirements, security threats, and existing solutions proposed for private targeted coupon delivery. Finally, we provided some discussion on future research directions.

References

[1] BEALES H. The value of behavioral targeting[EB/OL]. http://www.networkadvertising.org/pdfs/Beales_NAI_Study.pdf.

[2] PARTRIDGE K, PATHAK M A, UZUN E, et al. PiCoDa: privacy-preserving smart coupon delivery architecture[C]//The 5th Workshop on Hot Topics in Privacy Enhancing Technologies, Vigo, Spain, 2012: 94-108.

[3] RANE S, UZUN E. A fuzzy commitment approach to privacy preserving behavioral targeting[C]//The ACM MobiCom Workshop on Security and Privacy in Mobile Environments, Maui, USA, 2014: 31-35.

[4] BACKES M, KATE A, MAFFEI M, et al. ObliviAd: provably secure and practical online behavioral advertising[C]//IEEE Symposium on Security and Privacy, San Francisco, USA, 2012: 257-271.

[5] GUHA S, CHENG B, FRANCIS P. Privad: practical privacy in online advertising[C]//The 8th USENIX conference on Networked systems design and implementation, Boston, USA, 2011: 169-182.

[6] NATH S. MAdScope: characterizing mobile In-App targeted Ads[C]//The 13th Annual International Conference on Mobile Systems, Applications, and Services, Florence, Italy, 2015: 59-73.

[7] JUELS A. Targeted advertising and privacy Too[C]//Cryptographer's Track the RSA Conference, San Francisco, USA, 2001: 408-424.

[8] TOUBIANA V, NARAYANAN A, BONEH D, et al. Adnostic: privacy preserving targeted advertising[C]//The Network and Distributed System Security Symposium, San Diego, USA, 2010.

[9] HARDT M, NATH S. Privacy-aware personalization for mobile advertising[C]//The 19th ACM Conference on Computer and Communications Security, Raleigh, USA, 2012: 662-673.

[10] JIANG J, GUI X, SHI Z, et al. Towards secure and practical targeted mobile advertising[C]//The 11th International Conference on Mobile Ad-hoc and Sensor Networks, Shenzhen, 2015: 79-88.

[11] KHAN A J, JAYARAJAH K, HAN D, et al. CAMEO: a middleware for mobile advertisement delivery[C]//The 11th Annual International Conference on Mobile Systems, Applications, and Services, Taipei, China, 2013: 125-138.

[12] KODIALAM M S, LAKSHMAN T. V, MUKHERJEE S. Effective ad targeting with concealed profiles[C]//The 31st Annual IEEE International Conference on Computer Communications, Orlando, USA, 2012: 2237-2245.

[13] HUA J, TANG A, ZHONG S. Advertiser and publisher-centric privacy aware online behavioral advertising[C]//The 35th IEEE International Conference on Distributed Computing Systems, Columbus, USA, 2015: 298-307.

[14] CHEN R, REZNICHENKO A, FRANCIS P, et al. Towards statistical queries over distributed private user data[C]//The 9th USENIX Symposium on Networked Systems Design and Implementation, San Jose, USA, 2012: 169-182.

[15] AKKUS I E, CHEN R, HARDT M, et al. Non-tracking web analytics[C]//The ACM Conference on Computer and Communications Security, Raleigh, USA, 2012: 687-698.

[16] DWORK C. Differential privacy[C]//Proc. of the 33rd International Colloquium, Automata, Languages and Programming, Venice, Italy, 2006: 1-12.

[17] DWORK C, KENTHAPADI K, MCSHERRY F, et al. Our data,

ourselves: privacy via distributed noise generation[C]//Proc. of the 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, 2006: 486-503.

[18] RASTOGI V, NATH S. Differentially private aggregation of distributed time-series with transformation and encryption[C]//Proc. of the ACM International Conference on Management of Data, Indianapolis, USA, 2010: 735-746.

[19] SHI E, CHAN T H, RIEFFEL E G, et al. Privacy-preserving aggregation of time-series data[C]//The Network and Distributed System Security Symposium, USA, 2011.

[20] CHEN R, AKKUS I E, FRANCIS P. SplitX: high-performance private analytics[C]//The ACM Special Interest Group on Data Communication, Hong Kong, China, 2013: 315-326.

[21] NATH S, LIN F X, RAVINDRANATH L, et al. SmartAds: bringing contextual ads to mobile apps[C]//The 11th Annual International Conference on Mobile Systems, Applications, and Services, Taipei, China, 2013: 111-124.

[22] WU D, CHANG R K C. Analyzing android browser apps for file:// vulnerabilities[C]//The 17th International Conference Information Security, Hong Kong, China, 2014: 345-363.

[23] SON S, KIM D, SHMATIKOV V. What mobile Ads know about mobile users[C]//The 23rd Annual Network and Distributed System Security Symposium, San Diego, USA, 2016.

[24] LIU B, NATH S, GOVINDAN R, et al. DECAF: detecting and characterizing Ad fraud in mobile Apps[C]//The 11th USENIX Symposium on Networked Systems Design and Implementation, USA, 2014: 57-70.

[25] CRUSSELL J, STEVENS R, CHEN H. MAdFraud: investigating ad fraud in android applications[C]//The 12th Annual International Conference on Mobile Systems, Applications, and Services, Bretton Woods, USA, 2014: 123-134.

About the authors

JIANG Jinghua received the B.E. degree from the Xi’an Jiaotong University, Xi’an, China, in 2010. He is currently pursuing the Ph.D. degrees with the School of Electronic and Information Engineering, Xi’an Jiaotong University, and with the Department of Computer Science, City University of Hong Kong, Hong Kong, China. His research interests include cloud security, multimedia security, and mobile security. (Email: jjinghua2-c@my.cityu.edu.hk)

ZHENG Yifeng received the B.E. degree from the South China University of Technology, Guangzhou, China, in 2013. From September to December 2013, he studied at Zhejiang University, Hangzhou, China. He is currently pursuing the Ph.D. degree with the Department of Computer Science, City University of Hong Kong, Hong Kong, China. His research interests include cloud security, multimedia security, and encrypted data processing.(Email: yifeng.zheng@my.cityu.edu.hk)

SHI Zhenkui received the B.E. degree from the Southwest University of Science and Technology, Mianyang, China, in 2007, and the M.S. degree from the University of Electronic Science and Technology of China, Chengdu, China, in 2010. He is currently pursuing the Ph.D. degree with the Department of Computer Science, City University of Hong Kong, Hong Kong, China. His research interests include network security and mobile security. (Email: zhenkui.shi@my.cityu.edu.hk)

YAO Jing received the B.E. and M.E. degrees from the Shaanxi Normal University, Xi’an, China, in 2009 and 2012, respectively. She is currently pursuing the Ph.D. degrees with the School of Electronic and Information Engineering, Xi’an Jiaotong University, and with the Department of Computer Science, City University of Hong Kong, Hong Kong, China. Her research interests include cloud security and data privacy. (Email: jingyao7-c@my.cityu.edu.hk)

WANG Cong [corresponding author] received the B.E. and M.E. degrees from Wuhan University, Wuhan, China, in 2004 and 2007, respectively, and the Ph.D. degree from the Illinois Institute of Technology, Chicago, IL, USA, in 2012. He is currently an assistant professor with the Department of Computer Science, City University of Hong Kong, Hong Kong, China. His research interests include cloud and network security. (Email: congwang@cityu.edu.hk)

GUI Xiaolin [corresponding author] received the B.E., M.E., and Ph.D. degrees from Xi’an Jiaotong University, Xi’an, China, in 1988, 1993, and 2001, respectively. He is currently a professor and the deputy dean with the School of Electronic and Information Engineering, Xi’an Jiaotong University. His research interests include secure computation in open network systems, data privacy, and the Internet of Things. (Email: xlgui@mail.xjtu.edu.cn)