

Follow Your Silhouette: Identifying the Social Account of Website Visitors through User-Blocking Side Channel*

Takuya WATANABE^{†,††a)}, Nonmember, Eitaro SHIOJI[†], Mitsuaki AKIYAMA[†], Members, Keito SASAOKA^{††}, Nonmember, Takeshi YAGI^{†††}, and Tatsuya MORI^{††b)}, Members

SUMMARY This paper presents a practical side-channel attack that identifies the social web service account of a visitor to an attacker's website. Our attack leverages the widely adopted *user-blocking* mechanism, abusing its inherent property that certain pages return different web content depending on whether a user is blocked from another user. Our key insight is that an account prepared by an attacker can hold an *attacker-controllable* binary state of blocking/non-blocking with respect to an arbitrary user on the same service; provided that the user is logged in to the service, this state can be retrieved as one-bit data through the conventional cross-site timing attack when a user visits the attacker's website. We generalize and refer to such a property as *visibility control*, which we consider as the fundamental assumption of our attack. Building on this primitive, we show that an attacker with a set of controlled accounts can gain a complete and flexible control over the data leaked through the side channel. Using this mechanism, we show that it is possible to design and implement a robust, large-scale user identification attack on a wide variety of social web services. To verify the feasibility of our attack, we perform an extensive empirical study using 16 popular social web services and demonstrate that at least 12 of these are vulnerable to our attack. Vulnerable services include not only popular social networking sites such as Twitter and Facebook, but also other types of web services that provide social features, e.g., eBay and Xbox Live. We also demonstrate that the attack can achieve nearly 100% accuracy and can finish within a sufficiently short time in a practical setting. We discuss the fundamental principles, practical aspects, and limitations of the attack as well as possible defenses. We have successfully addressed this attack by collaborative working with service providers and browser vendors.

key words: web security and privacy, side channel, user identification

1. Introduction

The *Social web* has become ubiquitous in our daily lives. It includes not only popular social networking services such as Facebook and Twitter but also other forms of web services with social features, e.g., online services for video games such as Xbox Live and online auction/shopping sites such as eBay. Social web services facilitate interactions between people with similar interests. The widespread adoption of social webs has increased not only the number of users per

service but also the number of services used by each user. Mander [2] reports that Internet users have an average of almost seven social accounts.

Like many other web services, social webs have security and privacy concerns. What distinguishes social webs from other web services is that they have an intrinsic privacy risk; users are encouraged to share large amounts of personal/sensitive information on these services, e.g., personal photos, health information, home addresses, employment status, and sexual preferences. An attacker can collate various data from social web services to infer individuals' personal information. For example, as Minkus et al. [3] revealed, an attacker can recover a target's purchase history if s/he knows the target's eBay account. The purchases may include potentially sensitive items, e.g., gun-related items or medical tests. To protect privacy, an eBay user may use a pseudonym for his/her account name; even in such a case, however, an attacker who can link an eBay account with an account on Facebook, which encourages users to disclose their real name, can infer the identity of the actual person who purchased the sensitive items on eBay.

In this study, we introduce a side-channel attack that identifies the social account(s) of a website visitor. The key idea behind our approach is to leverage *user blocking*, which is an indispensable mechanism to thwart various types of harassment in social webs, e.g., trolling, unwanted sexual solicitation, or cyber bullying. Because user blocking is a generic function commonly adopted by a wide range of social web services, an attacker can target various social web services. In fact, our attack is applicable to at least the following various social web services: Ashley Madison, eBay, Facebook, Google+, Instagram, Medium, Pornhub, Roblox, Tumblr, Twitter, Xbox Live, and Xvideos. Because having an account with some of the services included on this list could involve privacy-sensitive information, any account identification can directly lead to privacy risks.

Our attack leverages the user-blocking mechanism as a means of generating the leaking signals used for the side-channel attack**. More specifically, we leverage the mechanism's inherent property that certain pages return different web content depending on whether or not a user is blocked from another user. Our key insight is that an account prepared by an attacker can hold an *attacker-controllable* bi-

Manuscript received March 10, 2019.

Manuscript revised July 26, 2019.

Manuscript publicized November 11, 2019.

[†]The authors are with NTT Secure Platform Laboratories, Musashino-shi, 180–8585 Japan.

^{††}The authors are with Waseda University, Tokyo, 169–8555 Japan.

^{†††}The author is with NTT Security (Japan) KK, Tokyo, 101–0021 Japan.

*Early version of this paper [1] was presented at IEEE EuroS&P 2018.

a) E-mail: takuya.watanabe.yf@hco.ntt.co.jp

b) E-mail: mori@nsl.cs.waseda.ac.jp

DOI: 10.1587/transinf.2019INP0012

**More precisely, our side-channel attack is classified as a *cross-site timing attack* that will be described in Sect. 3.1.

nary state of blocking/non-blocking, with respect to an arbitrary user on the service, and this state can be retrieved as one-bit data through cross-site request forgery and a timing side channel when a user visits the attacker’s website. We specifically refer to the property that enables this key action as *visibility control*, as an attacker is forcing another user to change how they see certain things in the system. Building on this primitive, we show that an attacker can use a set of controlled accounts to construct a *controllable* side channel, i.e., leaked data is completely under the attacker’s control. Using this mechanism, we show that it is possible to design and implement a robust, large-scale user identification attack mechanism on a wide variety of social web services. We note that the number of accounts required has a theoretically logarithmic relation to the number of users to be targeted, e.g., 20 attacker-prepared accounts are needed to cover 1 million users. The novelty of our attack is discussed further in Sect. 3.3.

We note that disabling our side channel, i.e., user blocking, requires careful assessment as it is a crucial function that is widely used on social webs. In Sect. 2.2, we discuss our analysis of the data for measuring the blocking behavior of more than 200,000 Twitter users [4] and revealed that 3,770 users have blocked more than 1,000 accounts. Our online survey also revealed that 52.3%/41.4% of Twitter/Facebook users have responded they have used the blocking mechanism before, and 92.4%/93.9% responded there should not be a limit on the number of blocks. These results suggest that neither disabling blocking nor posing a limit on it, is desirable from the viewpoints of the actual usage of the service and users’ expectations. Furthermore, as we show in Sect. 5.3, limiting the number of user blocks per account would not be an effective countermeasure owing to our additional technique, *user-space partitioning*.

To verify the feasibility of our attack, we performed extensive empirical studies using 16 existing social web services. We found that 12 of these services are vulnerable to the attack. Using 20 actual accounts, we found that the attack succeeds with nearly 100% accuracy under a practical setting.

Our contributions can be summarized as follows:

- We demonstrate that the *user-blocking* mechanism, which is an indispensable function widely adopted in various social web services, can be exploited as the leaking signals for a side-channel attack that identifies user accounts.
- In addition to the side-channel attack, we develop several techniques to accurately identify users’ accounts. We also reveal that this attack is applicable to many currently existing services. The attack has a high success rate of nearly 100%, and is high-speed, taking as short as 4–8 seconds in a preferable setting, or 20–98 seconds even in a crude environment with a large amount of delay.
- We discuss the principles, the practical aspects, and the limitations of this study, as well as some defenses against the attack.

- We have successfully addressed this attack by collaborative working with service providers and browser vendors.

2. Background: User Blocking

In this section, we first provide a technical overview of *user blocking*, which serves as a side channel used for the user identification attack. Next, we demonstrate that simply disabling/limiting this side channel is not a desirable solution against the attack from the viewpoints of actual usage and user expectations.

2.1 Technical Overview

User blocking is a means of blocking communication between two users. Note that some “blocking” mechanisms adopted by social web services are not *user blocking* per se but *message blocking*, e.g., “muting” or “ignoring”. While user blocking rejects a person access to your account, message blocking filters out all the messages (or notifications) originating from that person. Even if a person is blocked with message blocking, this does not necessarily mean that they do not have access to your online activities. In this paper, we will not focus on message blocking unless otherwise noted.

Social web services with *user-blocking* mechanisms have intrinsic web pages that change content depending on the status of the visitor, i.e., whether or not a visitor is blocked from another person. A typical example is a user profile data page, which provides information on a person such as a photograph (icon), a self-introduction, affiliation, recent posts/updates, etc. Figure 1 shows screenshots of some Facebook profile pages. In the non-blocked state, the user profile information is fully available; in the blocked state, these pieces of information are hidden. In addition to a user profile page, some social web services provide pages that reflect similar differences. A summary of such techniques is presented in Sect. 4.

To execute user blocking, a user typically clicks the “block” button set on the profile page of the person to be blocked or enters the account ID of the person in a text box shown on a dedicated page for user blocking. Even though official application programming interfaces (APIs) for performing user blocking are not necessarily provided on all social web services, to the best of our knowledge no services

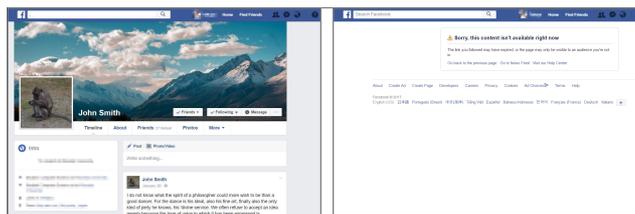


Fig. 1 The differences of appearance between non-blocking (left) and blocking (right) pages on Facebook.

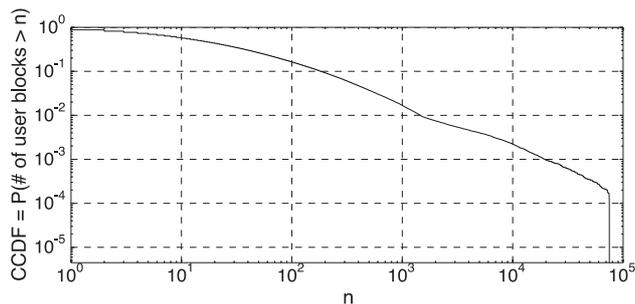


Fig. 2 Log-log CCDF of the number of user blocks per account on Twitter. Y-axis shows the fractions of accounts who are blocking n (X-axis) or more users. Mean value of the number of user blocks is 154.21.

adopt a special mechanism, such as CAPTCHA, to prevent automated user-blocking requests. Therefore, it is currently easy to perform the large-scale user blocking necessary to implement our user identification attack by using a script that emulates authentic requests or a headless browser.

2.2 Usage and Expectations

In this subsection, we discuss how many accounts do people block on social web services and why they do so. To answer the “how many” question, we first present statistics derived from the data collected at “Blocked By Me” [4], a web service that displays a list of users a person has blocked on Twitter[†]. The data, comprising the number of blocked users for 223,487 unique accounts, were collected from March 2011 to August 2017. As an individual may have used the web service for several times during the measurement period, we adopt the maximum value of the numbers of blocked users measured for each person. Figure 2 shows the log-log complementary cumulative distribution function (CCDF) of the number of blocked users per account. It is seen that the distribution is heavy-tailed, indicating that, although the majority of users blocked a small number of other accounts (median = 15), a non-negligible number of users had to block a large number of other accounts. For instance, 3,770 users blocked more than 1,000 accounts. Note that the rate-limit of access to Twitter API truncates the number of blocked users at 75,000; thus, users indicated in the figure as having blocked 75,000 users are likely to have actually blocked more. Besides this upper bound, there were several groups of accounts having the same large number of blocked accounts. They may be using a shared block list to evade various harassments. As checking the content of such lists is not feasible, some users may have simply cumulatively added new accounts to their block lists. These insights account for the reason why several users have a large number of blocked users.

Next, to answer the “why” question we recruited par-

Table 1 Demography of the expectations survey.

	# respondents	Gender	Age (Years)		
			10–29	30–49	50–
Facebook	198	F:54 M:46 (%)	31	60	9 (%)
Twitter	170	F:56 M:44 (%)	41	51	8 (%)

ticipants to take an online survey. As summarized in Table 1, the demography of the respondents shows that responses represent a diverse, cross-section of respondents. Key findings derived from the closed-ended questions are as follows: (1) 52.3%/41.4% of Twitter/Facebook users responded that they have used the blocking mechanism; (2) 92.4%/93.9% of Twitter/Facebook users responded that social web service should *not* limit the number of accounts a person can block on the service. This result indicates that users do not expect to have limitations on the number of blockable users. We also included the open-ended questions: “why do you block other users?” and “why do you think that there should be *no* limitation on the number of blocked users?” Typical answers to the first question include “do not want to read the unwanted messages/posts” and “not to be tracked by strangers/trolls/ex-friends/coworkers, etc.” Typical answers to the second question include “there are a huge number of spam/bogus accounts” and “just adding unwanted users to the blocklist is easy to maintain.”

The observations derived from the web service log analysis and the online survey imply that simply disabling our side channel, user-blocking, is *not* a desirable countermeasure against the threat from the viewpoints of actual usage of a service and users’ expectations.

3. Attack Overview

In this section, we give a brief overview of the attack. We present the threat model and the attack flow with a concrete example. We also elaborate on the novelty of the attack and how it compares to some of the existing works in this area.

3.1 Threat Model

In this attack, the attacker’s goal is to determine the social account of the visitors to her/his website. We present two possible attack scenarios under this goal. In the first, the attacker targets unspecified mass users in order to determine who visited the attacker’s website, for the purpose of, e.g., marketing. In the second scenario the attacker targets a limited number of users with already known identities, such as their names or email addresses, and wants to determine their anonymous accounts which cannot be searched for using such identities. In both scenarios, the visitor’s privacy is obviously breached, as the identity of the user or their private activity is revealed to the attacker without their consent. We further discuss the feasibility of building a target list in Sect. 7.

Our attack employs a *cross-site timing attack*, which is an attack that combines cross-site request forgery (CSRF) and a timing attack [5]. Cross-site timing attacks bypass the

[†]The dataset was provided courtesy of Gerry Mulvenna on August 14, 2017. Thus, there are negligible differences from the figures listed on his webpage archived on August 20 [4]. Note that the entire set was anonymized to protect user privacy.

same-origin policy and enable an attacker to obtain information using the target's view of another site, i.e., in our context, the attacker can know whether or not the target user is blocked by the attacker-prepared *signaling accounts*. As we detail in Sect. 4, the status of blocked/non-blocked can be estimated from the time a web server of a social web takes to load a web content, or the round-trip time (RTT), of the profile page of a signaling account. As such, we make the following assumptions, which we will discuss in additional detail in Sect. 7.

Attack Trigger. We assume that the attacker can somehow induce their target to visit a malicious website. For example, the attacker uses malvertising techniques [6] or simply sends out email messages, in which case they can also link e-mail addresses to social accounts. Further details on this are discussed in Sect. 7.

Log-in Status. We assume that a target person has logged into the social web services, i.e., that cookies are enabled on the person's web browser. This assumption plays a vital role in the success of the attack because the logged-in status triggers the difference between views of profiles of blocking and non-blocking accounts. Because the majority of web services, e.g., Facebook, have an automatic sign-in option, we consider this assumption to be reasonable.

User Device. We assume that the target person uses a PC when accessing the malicious website. This premise covers more than 70% of social web service users [7]. Users of mobile platforms typically access social web services through dedicated mobile apps instead of the web interface provided for mobile browsers. Therefore, we cannot easily apply the attack to a mobile device.

3.2 Attack Flow and Example

As illustrated in Fig. 3, our attack has two separate phases: the side-channel control phase and the side-channel retrieval phase. Below, we describe the steps in each phase with a concrete example. Note that some details are omitted for simplicity but will be described in later sections.

I. Side-Channel Control Phase

The purpose of the side-channel control phase is to construct user-identifiable side-channel data through user blocking. This phase is required just once before performing the attack.

Step 1. Target Enumeration: For a social web service of interest, the attacker first enumerates the users who will be the target of the attack. Let N be the number of targets. The attacker can target either mass (randomly sampled or even all) user accounts or a limited set of selected users (e.g., celebrities, high-level corporate officers) according to the attacker's purpose. Note that because this attack leverages CSRF, whether the account is closed (e.g., private, protected) is not a concern as long as the account is blockable. *In the example*, the attacker lists a small set of $N = 8$ users who will be the target of the attack. If the attack succeeds,

the attacker will be able to identify the accounts of these eight users whenever they visit the attacker's website while logged onto the social web services.

Step 2. Bit Assignment: The attacker prepares m accounts on the social web service where m is a number satisfying $2^m \geq N$; these accounts are referred to as "signaling accounts" and denoted as S_i , $i = 1 \dots m$. The attacker encodes a set of target users into bit arrays with length m , with the value of the i -th bit of each array corresponding to "block" (1) or "do not block" (0) by account S_i . The attacker can express a maximum of 2^m distinct target users, but at the cost of increase in m , the attacker can further add redundant bits to produce an error-correcting code.

In the example, the attacker prepares $m = 3$ signaling accounts, S_1 , S_2 , and S_3 , with each target user is mapped into distinct bit arrays of length m , as shown in the table. All possible bit patterns are mapped to the users and there are no redundant bits.

Step 3. Target Blocking: The attacker controls the signaling accounts to block each target user according to the bit array. Note that the number of blocking that must be performed per signaling account is approximately half of the total number of targets, as shown in the figure. It is not difficult to see that this requirement can be controlled at the cost of adding more redundant signaling accounts, i.e., the block/non-block table in the figure will become more sparse. *In the example*, S_1 is configured to block Erin, Frank, Grace, and Heidi, with the remaining four users left non-blocked (default). S_2 and S_3 are configured in a similar manner.

II. Side-Channel Retrieval Phase

The purpose of the side-channel retrieval phase is to identify the user utilizing the data retrieved through the timing side channel. This phase is executed every time a user accesses the attacker's website.

Step 1. User's Visit: When a user visits the web server under the control of the attacker, JavaScript code is downloaded and is executed on the user's browser.

Step 2. RTT Measurement: The JavaScript code (as detailed in Bortz [5]) measures the time taken to load the profile of the signaling accounts by sending HTTP requests to each of these accounts. Note that, as this is a CSRF, the request is issued on behalf of the user's account. Special RTT measurements are also performed to determine the threshold value used in the next step, but we omit the details here.

In the example, the script issues HTTP requests to the profile page of each of the signaling accounts — S_1 , S_2 , and S_3 — and receives the measurements of 214, 128, and 223 ms, respectively.

Step 3. User Identification: The attacker then tries to identify the user from the measurements acquired in the preceding step. Because the time needed to load the profile of a blocking account exhibits a statistical difference from that needed to load the profile of a non-blocking account, the sequence of measured time samples can be used to build a bit array of "blocked" and "non-blocked" states. Once the bit

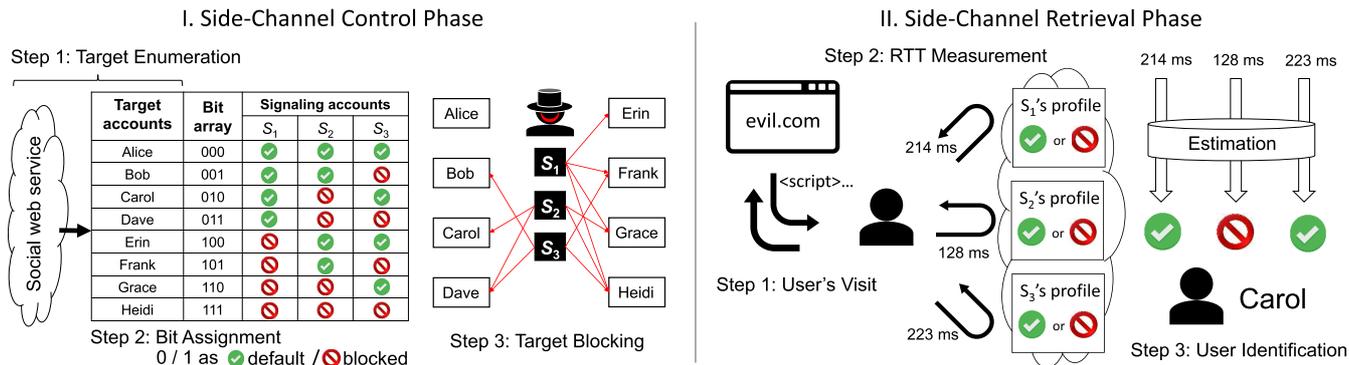


Fig. 3 Attack overview

array is recovered, the attacker does a lookup on the bit array map and identifies the user.

In the example, the measurements, 214, 128, 223 ms are compared against a threshold value of, say, 150 ms, and are determined to be non-blocked, blocked, and non-blocked, respectively. This result is represented as a bit array {010}, enabling the attacker to infer from the table that the user who visited the malicious site is Carol[†].

3.3 Novelty of the Attack

While our attack is certainly novel overall, its conceptual novelty lies primarily in the side-channel control phase rather than in the side-channel retrieval phase, which can be implemented using many different existing approaches in addition to that adopted in our implementation [5]. The side-channel control phase is made particularly novel by its use of the underlying concept of *visibility control*, which allows for the encoding and retrieving of arbitrary bits of data independent of what the side channel is. This flexibility inherently enables the attack to achieve account identification in a generic manner. By contrast, most similar methods that exploit browser side channels focus on stealing the content of a specific resource, limiting the acquirable data to that related to the targeted resource. Rather than studying such resource-specific side-channel acquisition methodologies, we questioned and exploited the design of general systems equipped with visibility-control features, e.g., user blocking. To the best of our knowledge, this concept has not been previously discussed in the literature despite its significant potential impact on nearly all major social web services currently operating.

We now compare our work to two of the major recent studies in this area. The goal of the first study was to retrieve various user data (e.g., age, contacts, search history) through several browser side-channel techniques [8]. The major difference between this work and ours is that it was somewhat focused on the development of individual techniques to acquire resource-specific side channels. Although this

[†]As we will detail in Sect. 5, when {000} is observed, it is still possible to distinguish Alice from non-target users by using two special accounts that do/don't block all the target users.

makes their methodology more powerful in the sense that it can even reveal a user's private information (e.g., search history), their methodology and goals were more service- and resource-specific. By contrast, the purpose of our work is to find user accounts and then link these with *all* available public information to which they are tied independent of the target resource used for sending side-channel data. Another similar study involved an attack based on browser history stealing [9], which, in the authors' words, shared a goal similar to ours of user identification or de-anonymization. This approach exploited the (now eliminated) mechanism allowing an attacker to infer a user's browser history to determine if the user belongs to certain groups based on the presence of access history to certain pages. Methodology-wise, the concept of repetitively identifying the groups to which a target user belongs, until to the point where the target can be uniquely identified, is conceptually similar to our approach. The main difference, however, is that our method allows for the construction of such groups in advance in an arbitrary manner. Thus, while our approach requires some initial setup effort, it has the advantage of being much more reliable in assuring identification (i.e., no ambiguity remains due to a lack of groups) as long as the side channel can be correctly retrieved.

4. User-Blocking Side Channel

This section aims to demonstrate that the differences between the time to load profile pages of blocked and non-blocked users can be used to perform a timing attack. In the following, we first look at the characteristics of the RTTs measured for blocked and non-blocked accounts. Next, we present several techniques that can increase the distinguishability of RTTs. Finally, after applying the RTT expansion techniques, we test whether the RTTs are statistically distinguishable using various social web services, which include popular social media such as Twitter and Facebook and other web services such as eBay and Xbox Live.

4.1 Characteristics of RTTs

Here, we briefly describe the setup for our experiments. We

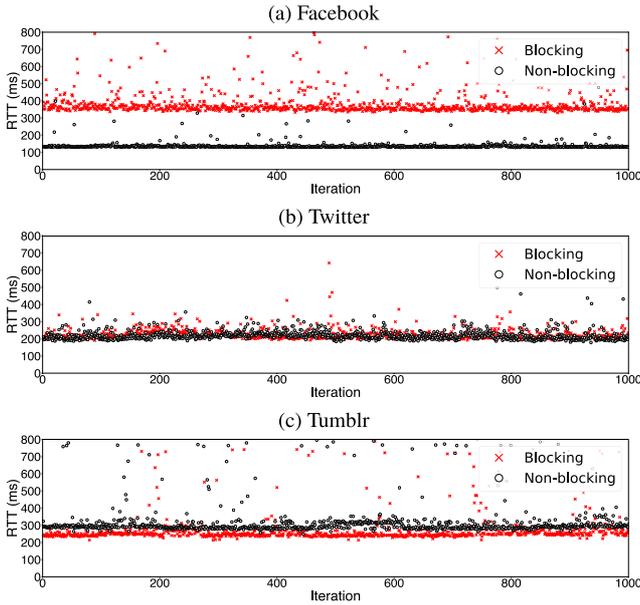


Fig. 4 Distributions of RTTs for blocking and non-blocking accounts.

executed a simple JavaScript code on a browser logged-in to a service with account A. The JavaScript issues GET requests to a page associated with an account which blocks A, and another page associated with an account that does not block A. Our objective is to see whether we can see the differences in the RTT measurements associated with these two types of accounts: blocking and non-blocking.

In the following, we characterize the measured RTTs using three social web services, Facebook, Twitter, and Tumblr as the representative examples. We study other services in the next subsection. Figure 4 shows the distributions of the measured RTTs[†]. For Facebook, there is a clear gap between the RTT distributions for blocking and non-blocking accounts. For Tumblr, even though two distributions are closer, we see the difference between the distributions. We study whether or not this slight differences can be used as the timing side channel in Sect. 6. For Twitter, the distributions suggest that there is no sufficient difference to distinguish their RTT difference. Nevertheless, we have discovered that it is possible to intentionally amplify their RTT difference by posting more content to the profile page. More details on this will be described in the next subsection.

Note that, while we see longer RTTs for non-blocking accounts on Tumblr, we see longer RTTs for blocking accounts on Facebook. It is natural that the profile pages of blocking accounts are loaded quickly because the content of these pages may be lighter than those of the profile pages of non-blocking accounts. While not conclusive, we conjecture that this could be because Facebook does not utilize its server-side on-memory cache at all when generating content

[†]The results of this section were measured in early 2017. As we will explain later in Sect. 8.3, this attack is no longer work for the services that have adopted countermeasures through cooperation with us.

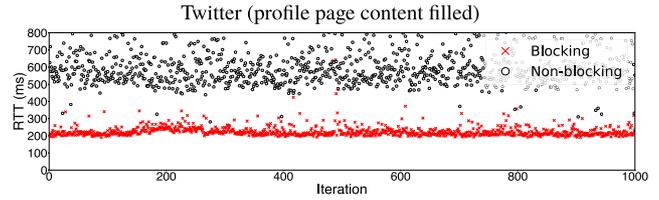


Fig. 5 Distributions of RTTs for blocking and non-blocking accounts, after filling the Twitter user profile page with content

for the case of blocked. In either case, we can distinguish between the blocked and non-blocked states using the RTT measurements.

4.2 Improving RTT Distinguishability

We present three techniques that can make the differences in the RTTs more prominent, i.e., these are the ways to make the timing attack more successful.

Change of content size. The first technique is to place as much information as possible on the user profile pages of the signaling accounts. This technique can increase the time to load the profile page when the signaling account of the page is visible to the target, i.e., the signaling account does not block the target. We performed a simple experiment using Twitter. We prepared two Twitter accounts, one with the default setting and another with the maximum amount of content (texts and URL links) that appears on the profile page. Figure 5 shows the RTT distributions after filling the profile page with large amounts of content. Comparing this with Fig. 4 (b) which shows the RTT distributions before adding the content, we now have a clear difference between blocked and non-blocked RTTs, suggesting that this technique can dramatically improve their distinguishability.

Use of different pages. Another technique is to make use of various pages other than the user profile page. In many cases, the page subject to blocking is the user profile page, which displays the user’s basic information or recent posts. However, depending on the service or their implementation, there are cases where observable differences do not appear on the profile page but do appear on other pages. For example, on eBay, a user cannot prohibit another user from accessing their profile page; however, a user can prohibit another user from bidding on the items they list. In other words, the content on the item page would yield a difference depending on whether the viewing user is blocked by the owner of the item. Leveraging this fact, by preparing an item beforehand and making the victim send requests to the item page instead of the profile page, the attacker would be able to observe the RTT difference required for the attack.

Similarly, Flickr does not prohibit a blocked user from viewing the blocker’s profile page, but it does prohibit the blocked user from sending a message to the blocker. More specifically, there is a page for sending a messages to other users and, if the sender is not blocked from the receiver, a text area and a submit button are displayed on the page;

however, if blocked, these objects are not shown and a warning message is displayed. Such a difference may also yield the RTT difference necessary for our attack.

In addition, some pages with AJAX-based implementation have a structure where after requesting and rendering the initial HTML content, they request additional content, e.g., a JSON content, from another URL using JavaScript's XMLHttpRequest. In some services, the blocked/non-blocked difference is only present in the JSON data that is acquired afterwards, instead of in the HTML content acquired first. The problem with this situation is that the RTT measurement script used for cross-site timing attacks does not actually render the acquired page content; therefore, the RTT of the content acquired afterward from JavaScript cannot be measured. In such cases, the attacker must directly send requests to the URL for the JSON data. In our investigation, we found that Tumblr and Xbox.com had this structure, but we were able to make the attack feasible by switching the request destination to the JSON URL instead of the HTML URL.

4.3 Distinguishability of RTTs

We tested whether the RTTs for blocking and non-blocking accounts were statistically distinguishable. To this end, we leveraged the Mann-Whitney U test, which is a nonparametric statistical test used to compare differences between two independent samples; it tests whether a randomly selected value from one sample is less than or greater than a randomly selected value from another sample. For our experiments, we picked 16 popular social web services. For each service, we measured the RTTs between blocking/non-blocking accounts and the blocked account. We applied the Mann-Whitney U test and computed the p -values. The results are summarized in Table 2. The results show that all services have low p -values and imply that the distributions are distinguishable in 12 out of 16 services when the significance level is 0.01.

Table 2 Social web services with user-blocking mechanism. $\Delta_{0.05}$ shows the difference in 5-percent tile values for blocked/non-blocked RTT measurements. Dist. is the distinguishability showing Y when the p -value less than 0.01. # of users are from various web resources as of May 2017

Service	Category	# users	$\Delta_{0.05}$	p -value	Dist.
Facebook	Social	1.96B	212 ms	<0.0001	Y
Instagram	Photo	700M	29 ms	<0.0001	Y
Tumblr	Microblog	550M	43 ms	<0.0001	Y
Google+	Social	540M	1,080 ms	<0.0001	Y
Twitter	Microblog	328M	312 ms	<0.0001	Y
eBay	Shopping	167M	589 ms	<0.0001	Y
PornHub	Porn	75M	9 ms	0.0034	Y
Medium	Forum	60M	332 ms	<0.0001	Y
Xbox Live	Game	52M	110 ms	<0.0001	Y
Ashley Madison	Dating	52M	8 ms	0.0097	Y
Roblox	Game	48M	98 ms	<0.0001	Y
Xvideos	Porn	47M	16 ms	<0.0001	Y
Quora	Forum	190M	5 ms	0.4561	N
Flickr	Photo	122M	1 ms	0.2678	N
DeviantArt	Art	65M	11 ms	0.0674	N
Meetup	Social	30M	9 ms	0.3878	N

5. User Identification Attack

In this section, we first formulate the user identification attack, which works on the basis of the two building blocks, user-blocking and cross-site timing attack. The attack introduces two functions, encoding and decoding, which are the functions an attacker can arbitrarily set to map target users and leaking information (RTTs). Next, we describe the techniques we developed for the timing attack. Finally, we present two extensions of the attack. These extensions aim to make the attack more successful.

5.1 Formulation

Let m and N denote the numbers of the signaling and target accounts, respectively. We configure m as the minimum integer value that satisfies $2^m \geq N$. If an attacker wants to target one million of accounts, m is configured to $m = 20$.

In the setup phase, an attacker creates a table that maps target user accounts to bit arrays with a length of m . Let U_i ($i = 1, \dots, N$) be the target user accounts. For each U_i , the table has a bit array entry, $B_i = \{b_1 b_2 \dots b_m\}$, where $b_j \in \{0, 1\}$ corresponds to a bit. We refer to the rule that maps U_i into B_i as *encoding*, i.e.,

$$B_i = \text{encode}(U_i).$$

Next, we configure the signaling accounts, S_j ($j = 1, \dots, m$) as follows. Let $\theta_{ij} \in \{0, 1\}$ ($i = 1, \dots, N, j = 1, \dots, m$) be an indicator function that satisfies

$$\theta_{ij} = \begin{cases} 1 & \text{if } b_{ij} = 1, \\ 0 & \text{else,} \end{cases}$$

where b_{ij} is the j -th bit of the bit array B_i . Then, for each signaling account, S_j , the account is configured to block the user U_i if $\theta_{ij} = 1$. Because each bit takes the value $b_{ij} = 1$ with a probability of 0.5, each signaling account needs to block approximately $N/2$ target accounts. One may instantly come up with a defense that poses a limit on the number of user-blocks an account can have. To thwart such a countermeasure, we propose a technique described in Sect. 5.3.

In the attack phase, the attacker sets up a malicious website and lets target users access it, following our threat model. As described in the previous section, using the timing attack, the website can secretly measure RTTs for the m of signaling accounts. Note that measurements can be parallelized to speed up the process. Let $\mathbf{R}_j = \{R_1, R_2, \dots\}$ be the sequence of RTT measurements obtained for the signaling account S_j . Using the techniques that will be described in the next subsection, we estimate whether or not the target user is blocked by S_j . Let $\widehat{b}_j \in \{0, 1\}$ denote the estimate of the blocked/non-blocked (1/0) from the RTT measurements, i.e.,

$$\widehat{b}_j = \text{est}(\mathbf{R}_j).$$

Using the entire estimates, we have the estimate of B as $\widehat{B} = \{\widehat{b}_1 \dots \widehat{b}_m\}$. Finally, we identify the target user using the table created in the setup phase; i.e.,

$$\widehat{U} = \text{decode}(\widehat{B}).$$

In the next subsection, the estimation, $\widehat{b}_j = \text{est}(\mathbf{R}_j)$, is described in detail.

5.2 Estimating Blocked/Non-Blocked Status

Prior to the actual attack, we determine whether or not a visitor of the website has been included in the target list, i.e., we employ a membership test. To this end, we prepare the following two reference accounts: a *closed account*, which blocks all users included in the list of target users, and an *open account*, which does not block any users at all. We first measure the RTTs for each of the closed and open accounts. The measurements consist of k_0 trials for each account, where we use multiple trials because the decision based on a one-shot measurement may have errors due to jitter in the RTTs. The idea is to compare the measured RTTs for closed/open accounts to see if they are significantly different. If we observe a significant difference, we can conclude that the visitor has been listed and continue the attack; otherwise, the visitor has not been listed and the attack procedure is terminated.

To determine if the measured RTTs are for the closed or open accounts, we again leverage the Mann-Whitney U test. Because the computation of the U test is simple and lightweight, the membership test can be completed immediately after we collect the RTTs. In this study, we adopted a significance level of $\alpha = 0.01$. We also need to configure the parameter k_0 . As shown in the next section, we empirically derived a conservative value of k_0 as $k_0 = 30$, which worked for various social web services.

After the attacker determines that the visitor is likely listed, the attacker moves to the next step. Let $C_{0.05}$ and $O_{0.05}$ be the 5th-percentiles of the RTT values measured for the closed and open accounts, respectively. We adopt the 5th-percentile as the threshold to eliminate outliers. Note that, even though we could use the minimum values for the RTTs as does the pathchar algorithm does [10], we observed that the RTTs could include small outliers, which could be caused by server-side mechanisms such as data caching or load balancing. These values are used as the thresholds to estimate the blocked / non-blocked state, i.e., for a measured RTT sequence for a signaling account S_j , we compute the 5th-percentile of \mathbf{R}_j as $R_{0.05j}$. We do not necessarily make k , the number of trials S_j , equal to k_0 . An attacker can adjust the k according to the his/her requirements for the trade-offs between accuracy and speed. If the obtained $R_{0.05j}$ is closer to $C_{0.05}$, the attacker estimates the visitor has been blocked by the signaling account S_j . Otherwise, s/he estimates the visitor has not been blocked by the signaling account; i.e.,

$$\widehat{b}_j = \begin{cases} 1 & \text{if } |R_{0.05j} - C_{0.05}| < |R_{0.05j} - O_{0.05}|, \\ 0 & \text{else.} \end{cases}$$

By continuing this process for all $j \in \{1, \dots, m\}$, the attacker can estimate the bit array of the visitor as $\widehat{B} = \{\widehat{b}_1 \dots \widehat{b}_m\}$. Finally, the bit array can be decoded into a user ID, $\widehat{U} = \text{decode}(\widehat{B})$, using the procedure we have shown in the previous subsection. Despite the simplicity of the procedure shown above, as we show later, it can estimate the closed/open states very accurately.

5.3 Extensions

Here, we introduce two extensions of the attack, *error-correction coding* and *user-space partitioning*, which aim to further improve the accuracy in noisy environments and to enhance the size of the target when the number of blocks per account is limited, respectively.

Error-correction Coding. Under a stable environment, accurately classifying a bit is not difficult since sufficient significant difference between blocked/non-blocked is present. This will also be shown later in Sect. 6. On the other hand, abnormal RTTs due to some irregular factors such as temporary server overload may lead to a bit-error. Needless to say, the infrastructures used in services such as those listed in Table 2 which host 30 million to 2 billion users tend to be quite resilient against such failures; nevertheless, we can still apply error-correction algorithm in order to eliminate even the slight possibility of identification failure due to noise.

In this paper, we adopt a Reed-Solomon code, which has a high error-correction capability and is relatively easy to implement. In fact, as we will demonstrate later, the use of the Reed-Solomon algorithm actually contributes to improving the estimation accuracy in a noisy environment. Note that other error-correction algorithms could be used for this purpose. To select the most suitable error-correction algorithm, one must take into account several factors such as the error probability distribution, the error characteristics such as bursts, and the requirements of the available computing resources. In this paper, we are focused on the proof of concept; therefore, we consider choosing the best error-correction algorithm to be out of the scope of this study.

The Reed-Solomon algorithm can correct up to $K/2$ symbol errors, where K is the number of redundant symbols and r (bits) is the size of the symbol. Because the number of bits initially allocated to each user is m , the number of signaling accounts that needs to be prepared by the attacker is $m + rK$, i.e., the attacker needs to prepare an additional rK extra signaling accounts. In the setup phase, the attacker first encodes the bit arrays allocated to each user using a Reed-Solomon encoder, and then blocks the accounts from the signaling accounts according to the bit values of the newly encoded bit array. In the attack phase, by decoding the bit arrays obtained via the cross-site timing attack using the Reed-Solomon decoder, the attacker can obtain an error-corrected bit array.

User-space Partitioning.

As described in Sect. 2.2, simply enforcing a limit on the number of blocks would violate a user's right to block

and may result in a serious degradation of the service quality. For services that still enforce a limit despite this negative impact, the technique shown below would be effective. Letting this limit to be L , the number of candidate target users covered for identification is also limited to L when using the procedures we have introduced up to this point. To lift this limitation, we can employ a technique we call *user-space partitioning*, which in this case splits candidate users into S user spaces each containing L users, allowing us to cover up to LS users in total.

In the setup phase, for each user space $j \in \{1, \dots, S\}$, an attacker prepares a reference account that blocks all users belonging to the j -th user space and the $\lceil \log_2 L \rceil$ of signaling accounts that are used to map the targets in the space. We also prepare the two reference accounts, the closed and open accounts, which are used as the basis of the RTT-based blocking/non-blocking estimation. In total, the number of signaling/reference accounts required is $S \lceil \log_2 L \rceil$ and $S + 1$, respectively.

In the attack phase, the attacker (1) identifies which user space the target user belongs to and then (2) identifies the target in the user space. In step (1), as in the procedures described in the previous subsection, for each of reference account, k requests are launched to determine the user space to which the target belongs. Note that the RTT values obtained here can be reused as the training data in step (2). In step (2), for each of the L users in the user space found in step (1), the same identification process is performed as explained earlier. Note that, because we use a different set of signaling accounts for each user space, the request URL for the cross-site timing attack must be changed depending on the outcome of step (1); however, this can be handled with conditional branches in the JavaScript code.

6. Field Experiments

In this section, we perform the field experiments. We first evaluate the key success factor of the attack – RTT measurement, which plays a vital role in classifying blocked/non-blocked status using the cross-site timing attack (Sect. 6.1). Next, we evaluate the feasibility of our user identification attack; namely, we study the identification success rate (Sect. 6.2) and time to complete the attack (Sect. 6.3).

6.1 Accuracy of Bit Array Estimation

Due to space and time constraints, we evaluated the accuracy using RTT values experimentally measured for the following three services: Facebook, Twitter, and Tumblr. As shown in Table 2, these services have the top number of users and, at the same time, had no limitations such as the limit on the number of blockable users at the time of the experiment. In addition, as mentioned in Sect. 4, each of these three services had different characteristics in the blocked/non-blocked RTT difference: relatively large, medium, and small, respectively.

The experiment was conducted by executing the

JavaScript on a consumer laptop PC and measuring the RTT. We argue that the user’s environment, i.e., network conditions and web browser engines, does not affect the success or failure of our attack. More precisely, our attack implementation absorbs the differences in the processing time caused by network jitters or the performance of the rendering engines. To demonstrate the assumption, we used the following three different network environments: wired LAN, Wi-Fi, and tethering. The wired LAN and Wi-Fi were connected to a commercial Internet service provider, and we assumed that this is the environment of PC users who are the main targets of our attack. To prove that our attack is feasible even in crude environmental conditions, we also used a tethering network hosted on an Android device connected to a 4G network provided by a mobile network carrier. In addition, we installed three of the most used browsers in the world, Google Chrome (v58), Mozilla Firefox (v53), and Microsoft Internet Explorer (v11). Note that trying all combinations was difficult due to space limitations. Unless otherwise noted, we used Google Chrome with a wired LAN.

Membership Test. We first tested the accuracy of the membership test. We measured the RTT for each of the closed and open accounts. As mentioned earlier, the measured RTT values are used for (1) the membership test and (2) deriving the thresholds for the bit classification, which will be described later. Note that an attacker needs to calibrate the thresholds before launching the attack because the RTT values depend on the geographical location and network environment.

We repeated the following experiment 100 times. While logged on to a target and non-target account, we launched k_0 trials for each account and decided whether or not the account was included on the list by applying the Mann-Whitney U test. We refer to the true positive rate (TPR) as the ratio of correctly deciding that the target was included in the target, and the true negative rate (TNR) as the ratio of correctly deciding that the target was *not* included on the target list.

Figure 6 shows the relationship between k_0 and TPR/TNR. When k_0 is small, we have a small number of samples to estimate the states. Nevertheless, thanks to the strong distinguishability of the RTT distributions, TNR was 0.97 for all k_0 , i.e., there were very few false negatives, which are events where the target account was estimated as *not* being listed. Second, for TPR, we saw degradation in the accuracy when k_0 was small, especially for Tumblr. As k_0 increases, however, the TPR approaches 1.0. When choosing the value of k_0 , it is preferable that the accuracy is consistent and that we see a sufficient difference in the samples. If k_0 is large, the accuracy will increase but the number of trials will also increase and the time needed for an attack would become too long. In this experiment, we empirically chose $k_0 = 30$, which achieved perfect estimations for all the services. We will use the values of $C_{0.05}$ and $O_{0.05}$ calculated from this k_0 as the thresholds used in the next step.

The measured RTT values can be affected by vari-

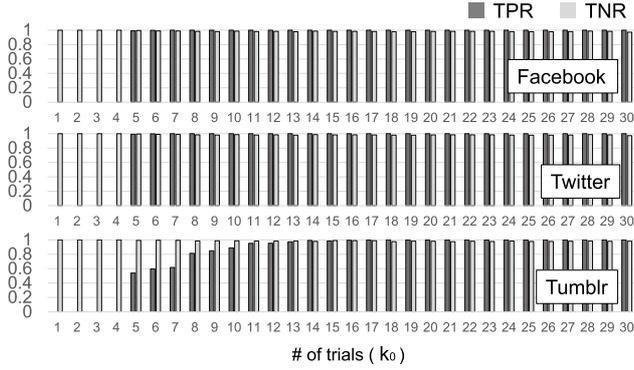


Fig. 6 Relationship between the number of trials (k_0), and TPR/TNR.

Table 3 TPR and TNR for under various conditions.

	Facebook		Twitter		Tumblr	
	TPR	TNR	TPR	TNR	TPR	TNR
Chrome/Wired	1.00	0.99	1.00	0.98	1.00	0.99
Wireless	1.00	0.98	1.00	0.98	1.00	0.99
Tethering	1.00	0.98	1.00	0.97	1.00	1.00
Firefox	1.00	0.98	1.00	1.00	1.00	1.00
IE	1.00	0.98	1.00	0.98	1.00	1.00

ous external factors such as network latencies or the type of browser. We studied how these factors affected the TPR/TNR. Table 3 shows the results. The number of trials was set to $k_0 = 30$.

Single Bit Classification. Next, we evaluated the accuracy of classifying a single bit into blocking or non-blocking. Again, we used three social web services, Facebook, Twitter, and Tumblr. For each service, we performed k trials of RTT measurements for each of two signal accounts with blocked/non-blocked states. We continued this process for 100 times and took the mean values of the following metrics. We refer to the true blocking rate (TBR)/true non-blocking rate (TNBR) as the rate of correctly detecting the blocking/non-blocking user as a blocking/non-blocking user, respectively. Table 4 shows the results. When $k \geq 20$, the detection becomes perfect for all the three services. Moreover, in a stable environment such as Facebook/Wired, the classification succeeds perfectly even with $k = 3$.

6.2 Attack Success Rate in the Wild

We now show the result of our experiment conducted in an environment imitating an actual attack scenario in the wild. We set the length of a bit array to $m = 24$, which can cover over 16 million users. In addition, we applied a Reed-Solomon code with a block length of 4 bits with eight redundant bits, which enables it to correct one block of error. According to the above setting, we prepared 34 accounts in total, which included 32 signaling accounts, a closed account, and an open account, with the appropriate blocking done against the users on the target list.

Regarding the targets, we assigned a random bit array of length 24 to each of the 10 social accounts we actually

Table 4 Accuracy of classifying a single bit for Wired (top), Wi-fi (middle), and Tethering (bottom)

k	Facebook		Twitter		Tumblr	
	TBR	TNBR	TBR	TNBR	TBR	TNBR
1	1.00	0.98	0.99	0.99	0.67	0.99
3	1.00	1.00	1.00	0.99	0.89	0.99
5	1.00	1.00	1.00	0.97	0.95	0.98
10	1.00	1.00	1.00	1.00	0.98	1.00
20	1.00	1.00	1.00	1.00	1.00	1.00
30	1.00	1.00	1.00	1.00	1.00	1.00
1	1.00	0.98	0.98	0.99	0.84	0.99
3	1.00	1.00	1.00	0.99	0.98	1.00
5	1.00	1.00	1.00	0.99	1.00	1.00
10	1.00	1.00	1.00	1.00	1.00	1.00
20	1.00	1.00	1.00	1.00	1.00	1.00
30	1.00	1.00	1.00	1.00	1.00	1.00
1	1.00	0.97	0.98	0.99	0.68	0.99
3	1.00	0.99	1.00	0.98	0.92	0.99
5	1.00	0.98	1.00	0.97	0.98	1.00
10	1.00	1.00	1.00	1.00	1.00	1.00
20	1.00	1.00	1.00	1.00	1.00	1.00
30	1.00	1.00	1.00	1.00	1.00	1.00

Table 5 Accuracy of the User Identification Attack.

	Facebook/wired	Twitter/WiFi	Tumblr/tethering
TNR	1.00 (20/20)	1.00 (20/20)	0.95 (19/20)
TPR	1.00 (20/20)	1.00 (20/20)	1.00 (20/20)
IDR	0.95 (19/20)	1.00 (20/20)	1.00 (20/20)
IDR/EC	1.00 (20/20)	1.00 (20/20)	1.00 (20/20)

own. We encoded these bit arrays using the Reed-Solomon code and calculated the bit arrays assigned as the redundant bits. We prepared 10 additional accounts which are not included in the list. For each of the 10 accounts on the target list and the 10 accounts on the non-target list, we logged in to and accessed the attacker’s website and evaluated if the account was correctly identified. We repeated the visit two times per account, resulting in a total of 40 identification trials.

As the parameters for the number of trials, we selected $k = 30$, which we experimentally determined yielded good accuracy. The service and network environment pairs we chose were Facebook/Wired LAN, Twitter/Wireless LAN, and Tumblr/Tethering. We refer to the TPR as the rate of correctly identifying a target to be included on the list, and the TNR as the rate of correctly identifying a non-target to be not included on the list. As mentioned above, we conducted the experiment twice with each account using 10 target accounts and 10 non-target accounts. Note that the denominators of TPR and TNR are 20. In addition, of the users who were identified as being included on the target list, we refer to the identified rate (IDR) as the rate of correctly identifying the user without the error-correction code, and refer to the identified rate with error correction (IDR/EC) as a similar figure but with error correction. In Table 5, we show the classification accuracy we obtained in this experiment.

The result shows that the experiment succeeded with extremely high accuracy. This was expected from the good results we obtained from the experiments in Sect. 6.1. For

Facebook/Wired, there was one failure case which identified the target as a wrong user. Examining the network log for this case revealed that some requests to one of the signaling accounts had returned 502 response code due to temporary server error. Our script measures the RTT even if an error code is returned, but since no content is returned, the response time would not likely be the one desired. This occurred with 3 of the requests over only 1 second of duration, but the RTT value had dropped to about 1/5 of the true RTT which was enough to cause a bit error. Nevertheless, applying the error-correction algorithm, we were successfully able to correct this bit which resulted in the success of identification. Note that, because we adopt the 5th-percentile, our attack is resilient to outliers which are too late, but it is prone to those which are too early.

Another case of failure was for Tumblr/Tethering, where a non-target user was incorrectly identified as a target. This is a rare case where a significant difference of around $p < 0.01$ happened to occur when comparing the two sets of 30 non-blocked requests. This example also benefited from the error-correction algorithm; without error-correction this visitor would have been identified as another user, but with Reed-Solomon code, although the error was not correctable due to too many errors, the error was still detectable. In such a case, we can still prevent mis-identification by concluding that the membership test failed and restarting the test.

6.3 Time to Complete the Attack

The shorter the time required for the attack, the more feasible the threat is. While the total number of requests can be calculated beforehand, the time required to complete these trials is dependent on the actual RTT; therefore, it needs to be evaluated experimentally. Figure 7 shows the relationship between the number of trials and the required time for each service.

The “upper bound” value shown for each service assumes the request with whichever has the larger of the blocked/non-blocked RTT values, that is, it assumes the case with the longest time needed for identification; i.e., it is the

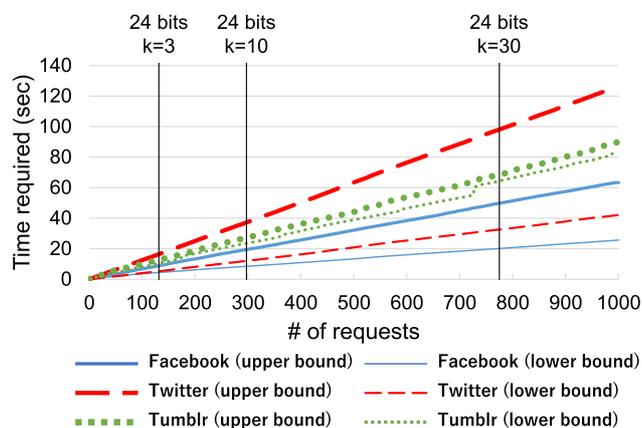


Fig. 7 Number of requests vs. time.

worst case. Conversely, the “lower bound” value assumes the request with whichever has smaller value of the two, that is, it assumes the case with the shortest time needed for identification; i.e., it is the most optimistic case. The number of trials issued in parallel was set to 6, which is the default maximum number of concurrent connections allowed on major browsers such as Chrome, IE, and Firefox.

The total number of requests needed to make an m -bits decision, or in other words, to identify the target within 2^m users, is $mk + 2 \times 30$ when $k_0 = 30$. For example, for $m = 24$, or targeting 16 million users, the total number of requests needed is 780 when $k = 30$. This would require 20–50 seconds for Facebook, 32–98 seconds for Twitter, and 64–68 seconds for Tumblr. According to Table 4, in the case of Twitter, we have sufficient accuracy even with $k = 10$. The number of necessary trials is 300 with this setting, and the time required is 12–37 seconds. Moreover, we can observe that we can achieve sufficient accuracy even with $k = 3$ on Facebook. The total number of requests is 132 which only takes 4–8 seconds.

7. Discussion

In this section, we discuss the attack’s principle, practical aspects, known limitations, and ethical considerations.

7.1 Principle of the Attack

We argue that the most fundamental assumption of our attack is the presence of the *visibility control* property in the system, that is, “given a multi-user web service, there exists a way for a (rogue) user to *control* what other users see, individually for each user”. To be more formal, the part “what other users see” can be replaced with “any observable side-effect of the system caused by a certain action taken by a user”. This assumption combined with a timing side-channel attack, which enables the attacker to steal this information from outside the system, is our attack’s big picture. Because closing a side channel completely is well-known to be difficult, we believe that this visibility-control assumption is the main principle of this attack. In the case of our scenario, the ability to build signaling accounts using user blocking corresponds to this principle.

We stress that other services under this assumption, even those without user blocking but with a similar mechanism such as group invitations or file access permissions, may also be subject to a similar class of attack. Still, the social web/user-blocking example that we used in this paper is by far the most practical application. This is likely because it satisfies several additional conditions: (1) the control can be done without the target’s approval or notification and (2) the control can be done at a fine granularity, i.e., the different bits of information assignable per user is large. More specifically in our case, condition (1) is almost always achieved as an inherent nature of user blocking and condition (2) is achieved with unlimited granularity, in theory, via the creation of an arbitrary number of signaling accounts. Even

though we omit further discussions concerning the presence of other such properties or the exploitability of similar systems, we believe that there is a need for further study concerning this subject.

7.2 Practical Aspects

Here we describe some of the applications and characteristics that extend and strengthen our attack from a practical perspective.

Target List Building. With our attack, it is assumed that an attacker knows the identity of a victim account in advance, implying she/he has a list of identities for the target users. Thus, building a target list is a crucial factor for attack success. A *limited set of selected accounts* and *unspecified large number of accounts* are both attractive target candidates for an attacker. The former will include famous people (e.g., politicians, corporate officers, celebrities, YouTubers) or intimates of the attacker. In such cases, the target list can be built based on the attacker's knowledge or the existing open list [11]. The latter[†] is realistic by using automation techniques, in particular, chasing the chain of friends' friends [12], [13]. Ugander et al. [14] revealed that 99.91% of Facebook users belong to a single huge relationship graph, and Dey et al. [12] found that 82.73% of users publish their friend lists. Depending on the service feature, user searches and user groups may be useful for user enumeration. Some services, such as Twitter, can list all accounts simply by incrementing a numerical user identifier.

Identity Linking. User identification is only threatening if the identity is linked with another piece of information meaningful to the attacker. In the case of our attack, the most basic form of linking can take the form of linking the user's identity with *the fact that the user has visited the website* prepared for the attack. In this case, if the web content reflects the visitor's preference in any way, it may become a privacy concern. This is suitable not only for advertisement or access analyses, but also for various social engineering attacks or for blackmailing those who have accessed sites hosting pornographic content or illegal content such as pirated software. In addition, our attack can be implemented to reveal accounts on multiple services simultaneously and linking these accounts together could significantly worsen the impact of a privacy leak.

Another form of linking occurs when a person is induced to access the web server via an extra hop through another medium, resulting in a linking between the target's identity and the medium used. For example, on a social web service where the target's identity is already known, an attacker can send the target a message containing the URL of the web server. Note that this would allow the attacker to link even the web services which our attack cannot be ap-

plied to. Similarly, we can link non-web services, such as email or mobile text messages, which would result in linking an email address or phone number with a social account. Further, we can also link the target's *physical identity*, such as the target's physical presence or their residence, by placing or mailing a physical object, e.g., a poster or a flier, with URL, QR code or NFC tags printed on them. Note that, even though it may appear that revealing additional identities of a target when the target's other identities are already known is not so significant, it could lead to the identification of a target's anonymous account that cannot otherwise be discovered in a straightforward way.

Group Identification. Even though we have focused on the goal of *user* identification in this paper, we can easily extend this goal to *group* identification, that is, identifying not the user's exact identity but more general properties such as gender, nationality, or interests. The attacker could map each user to a bit array corresponding to the target attribute collected from the structured information available on the social web service. Note that this can be seen as a generalization of the user-space partitioning described previously, where a user space corresponds to a group of users with an arbitrary size mapped to a certain attribute. Group identification can be used by advertisement providers to track the visitor's attributes without unnecessarily revealing their user account. Note that the number of bits required for group identification would typically be much lower than that for user identification, making this attack significantly easier to execute than user identification.

Authentication-backed Identification. One major strength of our approach is that it is backed by the identity information guaranteed by the authentication system of the service, making it resilient against spoofing or misidentification, both of which many other methodologies suffer from. To give a simple example, when using an IP address for identification or even tracking, IP spoofing or ambiguity due to NAT or dynamic IP would interfere with this process. Note that social web service accounts are increasingly used as a building block in the modern web's authentication infrastructure. It is still possible to perform spoofing and one way is to create an account trying to mimic one's identity; however, scrutinizing the account content would usually easily reveal whether it is a spoofed account. Another way is to use a stolen account, but in this case, the victim user should be worried about much more serious problems than privacy leakage. In addition, because authentication is independent of the environment, it enables cross-environment (e.g., cross-device and cross-browser) identification and tracking, which is often difficult to achieve using other approaches.

7.3 Limitations

Login State Persistence. Our attack relies heavily on the assumption that the target user's service login state is alive while the user browses other websites. This assumption

[†]Even in the latter case, it is reasonable for the attacker to create subsets of all enumerated users according to user attributes (nationality, age, school, company) and reduce the number of user blocks. We discuss the limit of user blocking in Sect. 7.3.

is reliant on the web cookie mechanism; therefore, the cookie's expiration time or the user configuring the browser to clear cookies on closing the browser may affect the availability of our attack. Social web services, fortunately, tend to set a relatively long or even no expiration time, as seen in the commonly available "keep me logged in" features [15]. This is likely due to the incentives to service providers from a marketing perspective, e.g., tracking and advertisement, contrary to security-critical services such as Internet banking that set a short expiration time. In addition, users would lose the convenience of being able to access the service without the need to login every time, which may be a disappointing trade-off, especially for social web services which often assumes constant usage. Note that, simply determining whether a user is logged in to certain services can be accomplished in much more lightweight ways [16], which can also be used in our attack to pre-select the services to be targeted.

User Environment. Our results from field experiments in Sect. 6 indicate that our attack is almost completely successful with various network environments and browser engines. However, we recognize that we could not consider all Internet user conditions. Especially, a non-negligible portion of users today access social web services from their mobile devices, so whether or not the attack is feasible in this realm is an important question to explore. For recent mobile platforms such as Android and iOS, the mechanics of most web browsers as well as the effective performance of the hardware and network are not significantly different from those of a PC; therefore, they are expected to yield sufficient RTT differences making our attack feasible. We partially proved this in our experiment with the tethering environment. The primary concern instead is the unique software ecosystem of mobile devices: many services encourage users to use a service-dedicated app instead of a browser to access their service. Even though some collaborative features such as social plug-ins or single sign-on may still urge some mobile users to log on via a browser, this ecosystem will surely limit the target coverage of our attack to a certain degree. We believe that a possible attack vector for this scenario which may need an attention might be an exploitation of a mobile platform-specific side channel, e.g., Android's Intent and shared memory [17], to bypass the app sandbox, analogical to how our attack exploited a browser timing side channel to bypass the same-origin policy, but we leave further discussions on this for a future study.

Limits on Blocking. For most services, limitations on the total number of users allowed to be blocked or the rate at which blocking requests can be issued from a single account are not explicitly stated. We have experimentally confirmed that at least ten million users on Twitter and three million users on Facebook and Tumblr were actually blockable over five days using a single account, and only DeviantArt and eBay seems to have had a limit on the maximum number of blocks per account. Also, Instagram appears to have had a limitation on the rate, i.e, the number of accounts that can be

blocked per minute. As we have shown in Sect. 2.2, neither disabling blocking nor posing a limit on it, is desirable from the viewpoints of the actual usage of the service and users' expectations. However, having limits on the total number of users to be blocked blocking may interfere with the process of building a high-coverage signaling account. Still, user-space partitioning would help alleviate this limitation and much of the effort for building signaling accounts is required just once, implying that attackers are not so exceedingly time-constrained when performing this task.

Length of Visit. As shown in Sect. 6, the attack can be executed in a realistically short time. In certain circumstances, however, such as when the RTT is high or when there is a need to use user-space partitioning, which increases the number of requests, it may be difficult to keep the user on the same webpage long enough for the JavaScript code to finish. Even if the attack duration is short, because the behavior of a user is often unpredictable, a shorter attack is always preferable. A trivial approach to this problem is to prepare webpage content that is sufficiently "attractive" to cause the users to stay longer, but this is very user specific. Another solution is to save and restore the attack state between multiple attack sessions. By having the JavaScript code send partial results to the server as it attacks, even if the attack terminates before finishing, the attack can be resumed at another session from where it left off. Training data may be reused or not depending on the "distance" between each attack session, e.g., the time elapsed between sessions. Another solution is to open pop-up windows in the background or a tab and execute the attack there, hoping that the user would not notice or care to close it immediately.

7.4 Research Ethics

In Sect. 2.2, all the data have been collected with user consent, and we followed guidelines presented by the ethics committee of Waseda University. To evaluate the feasibility and impact of the attack techniques on social web service users, experimenting with attacks on actual social web services cannot be avoided. All attacks in our experiment were checked manually and only generated a restricted amount of request. As a result, our experiment was carefully controlled and only generated a restricted amount of traffic (requests), which did not increase the workload of the sites and did not undermine the quality of their services. Furthermore, our experiment performed against our own accounts. Therefore, actual users of the services we examined were not directly involved in our attacks.

8. Defense

In this section, we discuss defensive measures that can be taken against our attack. We also present several countermeasures adopted by today's popular social web services. These countermeasures were developed with our aid.

8.1 Server-Side Defenses

Access Validation. Token-based defenses are widely adopted to prevent CSRF attacks in general. The server appends a one-time random string, or token, to each URL link generated and verifies it when the link is accessed. A defense mechanism that validates a referrer is also effective because it can accept requests from the whitelisted URLs while rejecting all other illicit requests. These prevent any third-party from generating a valid link; therefore, the attacker will not be able to receive valid responses containing information useful for the attack as long as the access validation process is applied before the block checking at the server side. A major drawback of these defenses is that legitimate requests are also affected and result in consequences such as breaking search engine results or prohibiting any means of link sharing, including those on blog posts and emails. Promising approaches which acquire user contents by using JavaScript's XMLHttpRequest with a valid token such as *placeholder* [8] and *double-submit cookie* [18] have been proposed, but they still require a change in the system architecture design and also the delay caused by the extra hop may negatively affect the user experience.

Response Time Control. The server could adjust the response time to minimize the block/non-block RTT difference. One approach is to artificially equalize the response times by adding delays to whichever has the shorter response time. Another approach is to randomize the response time by injecting delays of random lengths. However, either approach would impose a non-negligible performance degradation experienced by the user. In general, this type of timing side-channel defense is difficult to perfect; the profound study results in this area provide advanced attackers with various ways to amplify such differences at the cost of some increased effort, as we also have exemplified in this paper. In addition, the network delay is often uncontrollable from the service side so a perfect control is difficult to attain from the server side. Note that such types of server-side defenses are often thwarted by other timing side-channel approaches, such as those leveraging the content cache [8].

Usage Restriction. Our attack, when implemented in a straightforward manner, may exhibit behavioral characteristics not usually seen in the normal usage of the service. One case of such an anomaly would occur in the preparation process of a signaling account, which requires a massive number of blocking requests to be issued within a short time. Another is in the process of launching the attack from a browser, which causes an abnormal number of GET requests to be issued. The service can either restrict this in the form of the rate limit, CAPTCHA, or some means of heuristic anomaly detection. However, these defenses are expected to function only as a mild mitigation, because advanced attackers have historically been able to circumvent these types of defenses. The most extreme form of restriction is to remove the user-blocking capability from the ser-

vice. All these types of restriction-based measures, however, lead to an undermining of the ability to suppress those who truly needs to be blocked, which may result in a degradation of the service quality.

8.2 Client-Side Defenses

User. Defenses that can be taken by a user alone are limited to quite trivial ones. One approach is to isolate the browsing environment in which the web service is used, from that used for other purposes. This can be done, for example, by using the private browsing feature commonly available in modern browsers, logging out of the service when not in use, or simply using a different browser. Another approach is to restrict the execution of JavaScript using browser plug-ins such as NoScript [19], which would severely impair the attacker's capability to carry out such an attack. Obviously, all of these measures greatly increase the user's cost of not only using the service but also web browsing in general. Further, it would deactivate some features such as social plug-ins or advertisements that benefit both of the user and the service provider.

Web Browser.

SameSite is a cookie attribute that allows flexible control of sending cookies in cross-site requests. It is necessary for a browser to adopt this feature, and the web service explicitly declares it in the HTTP header. Since browsers behave as if a user is not logging in when they make cross-site requests from a third-party site to a social web service, the difference in RTTs between blocking and non-blocking states disappears in principle. A Web service that adopts *SameSite* has two options: *samesite=strict* and *samesite=lax*. For the *strict* option, browsers remove cookies from any cross-site request, including redirects and link clicks. As with access validation, legitimate page transitions may be disturbed as well. For the *lax* option, however, browsers do not prevent the sending of cookies for cross-site requests with top-level navigation. Thus, the *lax* option can be used to defend against various CSRFs including timing attacks while not sacrificing user experience. However, *samesite=lax* still interferes with the functionality of some social plug-ins that leverage cookies from the social web such as Facebook Comments.

Equalizing the response times, for example, by injecting delays to the processing time, is also a possible measure that can be taken on the browser side. Further, the detection of anomalies such as frequent errors resulting from failed rendering may be another option. However, these approaches are often only viable for a certain class of timing side channels; they tend to be thwarted eventually by other newly developed timing attacks using different approaches, as exemplified by the attack using the browser cache mentioned in another study [8]. Cao et al. [20] proposed *Deterministic Browser*, which prevents JavaScript program from accessing the physical clock when a secret event is running at the same time. Instead, they provide an auxiliary clock that indicates virtual time to make timing attacks impossi-

ble. This approach works against our attack, but it can hinder the browser rendering for some websites.

8.3 Responsible Disclosure

Even though the attack technique in this paper does not arise from a specific social web service, according to the principle of responsible disclosure, we have reported the details of our attacks and the experimental results to the relevant social web service providers and browser vendors to mitigate the attacks and improve future security design of social web.

Through these activities, Twitter has been able to prevent the threat of account identification by changing their specifications [21] to improve their security mechanism. Their primary defense approach is using the *SameSite* attribute. This is currently working for most users, as we and Twitter encourage major browser vendors (including Microsoft and Mozilla) to adopt *SameSite*. For older browser versions, Twitter's other mitigation still prevents our attack. It is referrer-based access validation, but a user with invalid referrers is not simply denied. Instead, such a user is redirected to a stepping-stone page with no difference in RTT with only minimal JavaScript code then redirected again to the destination page via JavaScript. This works as a defense because the time taken to redirect via JavaScript cannot be measured under the constraints of the same-origin policy. While this approach slightly increases the number of requests, users coming from external links reach the desired content on the social web. However, the user cannot use browsers (or browser extensions) that remove the referrer for privacy.

Although we refrain from disclosing the brand names, several other service providers and browser vendors have already finished implementing defenses and some are in the process.

9. Related Work

We present previous studies concerning timing attacks, which is the fundamental technique of our method uses to compromise user's privacy. In addition, we introduce other side-channel leaks based on the browser functionality and methods to identify and track users.

9.1 Web-Based Timing Attacks

A timing attack is one type of side-channel attack that has been studied primarily in cryptography for more than two decades. It typically exploits the execution time or power consumption of a cryptosystem to infer secret key and private information [22], [23]. Studies of timing attacks have expanded to web-based systems regardless of the cryptosystem that exploits the communication time and size of the web content. Bortz et al. presented a pioneer work on web-based timing attacks; they classified web-based timing attacks into *direct timing* and *cross-site timing* [5]. Our proposed method is classified as a web-based cross-site timing

attack.

A direct timing attack directly measures the response times from a system, e.g., a website, to extract private information from a system. Bortz et al. proposed a method to expose valid user names and the number of private photos from a website by measuring the response time of HTTP [5].

Cross-site timing attacks indirectly measure the response times or content size of web on a browser to extract private information from a browser or website. It enables a malicious website to obtain information about the target browser's view of another website using cross-site content that often violates the same-origin policy [24]. Methods to break the same-origin policy and their countermeasures have been presented since 2000 [25]–[28]; however, the many of cross-origin techniques are still effective on modern web browsers. Liang et al. leveraged several CSS features to indirectly monitor the rendering of a target resource [29]. Goethem et al. proposed a cache-based timing attack using HTML5 functionalities, which can bypass the same-origin policy, to estimate the size of a cross-origin resource [8]. Gelernter et al. presented a *cross-site search attack* on well-known web services to distinguish between the loading time of empty and full responses, which enables an attacker to distinguish sensitive data of target users in the records of the web services [30]. Jia et al. demonstrated a *geo-location inference attack* on well-known web services, by using the load time of location-sensitive resources left by geography-specific websites (e.g., Google's local domain) [31]. Our method is not new in the context of cross-site attacks; however, the idea is unique in that user blocking, which is a fundamental functionality of social webs, can be used to distinguish between the blocked and non-blocked states, consequently, to identify their social accounts.

9.2 Side-Channel Leaks on Browsers

A side-channel attack on a browser without timing features is another class of privacy attack. To infer the status of a cross-origin resource, Lee et al. developed a *URL status identification attack* using `ApplicationCache` that exploits cross-origin resource caching [16] and they suggested advanced privacy threats using this attack, e.g., login status determination and internal web server probing. A *history-stealing attack* is a typical attack that extracts the browsing history of URLs [9], [32]. This attack depends on the fact that a web browser handles CSS properties of URL hyperlinks differently depending on whether the URL was previously accessed by the web browser [33], which leads to allowing a client-side script to access such properties. To fix this, Baron proposed a solution that blocks scripts from accessing the CSS properties of hyperlinks, and all popular browsers (e.g., Firefox, Chrome, Safari, and IE) have adopted this solution. As a result, this type of history stealing attack no longer works in the latest versions of these browsers [34], [35].

9.3 Social Account Identification

While various methods have been proposed to effectively track browsers on the Internet (e.g., cookies, browser cache, and browser fingerprints [36]–[38]), these tracking methods focus on identifying distinct browsers rather than the user of the browsers. The goal of our proposed method is to identify the user (i.e., the social account) which differs from the above browser tracking methods. Many of the studies introduced in Sects. 9.1 and 9.2 mentioned that their proposed methods could be used for inferring the status of social account or identifying social account [5], [9], [16]. The difference of response time of login page was used for inferring account validity [5]. With a similar motivation, conditional redirections of the HTTP URLs was used for distinguishing whether a victim web browser is logged in to the web service [16]. The combination of *group membership information*, e.g., group ID or group directory in browser's access history, was used for identifying a social account [9]. These differences are extracted from previously provided pages, e.g., login pages and group membership pages. In contrast, our method is unique in that an attacker can fully control the visibility of pages in order to create discriminable differences.

10. Conclusion

This work presents a practical side-channel attack that identifies the social account of a user visiting the attacker's website. It exploits the user-blocking mechanism, or the visibility control property, commonly available in most social web services today to create a controllable side channel that provides the attacker with complete and flexible control over the leaked information, be it informative enough to uniquely identify the user or be it highly resilient to noise. With experiments, we demonstrated that our attack is in fact applicable to current mainstream social web services today and we argued that defending against this threat would not be easy without imposing a negative impact on the relevant services. It is ironic that the blocking feature designed to suppress harmful users can now be turned against harmless users; some form of mitigation is urgent and a reworking of the design of this feature is suggested and major services and browsers adopted new security features.

References

- [1] T. Watanabe, E. Shioji, M. Akiyama, K. Sasaoka, T. Yagi, and T. Mori, "User Blocking Considered Harmful? An Attacker-Controlled Side Channel to Identify Social Accounts," *Proc. of IEEE EuroS&P*, 2018.
- [2] J. Mander, "Internet Users Have Average of 7 Social Accounts," <https://blog.globalwebindex.com/chart-of-the-day/internet-users-have-average-of-7-social-accounts/>, 2016.
- [3] T. Minkus and K.W. Ross, "I Know What You're Buying: Privacy Breaches on eBay," *Proc. of PETS*, 2014.
- [4] G. Mulvenna, "Blocked By Me," <https://web.archive.org/web/20170820142006/blockedby.me/>.
- [5] A. Bortz, D. Boneh, and P. Nandy, "Exposing Private Information by Timing Web Applications," *Proc. of WWW*, 2007.
- [6] Z. Li, K. Zhang, Y. Xie, F. Yu, and X. Wang, "Knowing Your Enemy: Understanding and Detecting Malicious Web Advertising," *Proc. of ACM CCS*, 2012.
- [7] K. Young, "GlobalWebIndex Blog," <https://blog.globalwebindex.net/chart-of-the-day/8-in-10-social-networking-on-mobile/>, 10 2016.
- [8] T.V. Goethem, W. Joosen, and N. Nikiforakis, "The Clock is Still Ticking: Timing Attacks in the Modern Web," *Proc. of ACM CCS*, 2015.
- [9] G. Wondracek, T. Holz, E. Kirda, and C. Kruegel, "A Practical Attack to De-anonymize Social Network Users," *Proc. of IEEE S&P*, 2010.
- [10] V. Jacobson, "Pathchar," <https://www.caida.org/tools/utilities/others/pathchar/>, 1997.
- [11] "Fan Page List," <https://fanpagelist.com>.
- [12] R. Dey, C. Tang, K. Ross, and N. Saxena, "Estimating Age Privacy Leakage in Online Social Networks," *Proc. of INFOCOM*, 2012.
- [13] B. Krishnamurthy, P. Gill, and M. Arlitt, "A Few Chirps about Twitter," *Proc. of Workshop on Online Social Networks*, 2008.
- [14] J. Ugander, B. Karrer, L. Backstrom, and C. Marlow, "The anatomy of the facebook social graph," *arXiv preprint arXiv:1111.4503*, 2011.
- [15] G. Kontaxis, M. Polychronakis, A.D. Keromytis, and E.P. Markatos, "Privacy-Preserving Social Plugins," *Proc. of USENIX Security*, 2012.
- [16] S. Lee, H. Kim, and J. Kim, "Identifying Cross-origin Resource Status Using Application Cache," *Proc. of NDSS*, 2015.
- [17] C.Q. Alfred, Q. Zhiyun, and M.Z. Morley, "Peeking into your app without actually seeing it: Ui state inference and novel android attacks," *Proc. of USENIX Security*, 2014.
- [18] OWASP, "Cross-Site Request Forgery (CSRF) Prevention Cheat Sheet," [https://www.owasp.org/index.php/Cross-Site_Request_Forgery_\(CSRF\)_Prevention_Cheat_Sheet](https://www.owasp.org/index.php/Cross-Site_Request_Forgery_(CSRF)_Prevention_Cheat_Sheet).
- [19] G. Maone, "NoScript Security Suite," <https://addons.mozilla.org/en-US/firefox/addon/noscript/>, 2016.
- [20] Y. Cao, Z. Chen, S. Li, and S. Wu, "Deterministic Browser," *Proc. of ACM CCS*, 2017.
- [21] K. Kufluk and G. Baker, "Protecting user identity against Silhouette," https://blog.twitter.com/engineering/en_us/topics/insights/2018/twitter_silhouette.html.
- [22] P.C. Kocher, "Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems," *Proc. of CRYPTO*, 1996.
- [23] P. Kocher, J. Jaffe, and B. Jun, "Differential Power Analysis," *Proc. of CRYPTO*, 1999.
- [24] J. Ruderman, "The same origin policy," <http://www.mozilla.org/projects/security/components/same-origin.html>, 2001.
- [25] CERT, "Malicious HTML Tags Embedded in Client Web Requests," <https://www.cert.org/historical/advisories/CA-2000-02.cfm?>, 2 2000.
- [26] E. Kirda, C. Kruegel, G. Vigna, and N. Jovanovic, "Noxes: A Client-Side Solution for Mitigating Cross-Site Scripting Attacks," *Proc. of ACM Symposium on Applied Computing*, 2006.
- [27] P. Vogt, E. Kirda, C. Kruegel, and G. Vigna, "Cross-Site Scripting Prevention with Dynamic Data Tainting and Static Analysis," *Proc. of NDSS*, 2007.
- [28] B. Stock, S. Lekies, T. Mueller, P. Spiegel, and M. Johns, "Precise Client-side Protection against DOM-based Cross-Site Scripting," *Proc. of USENIX Security*, 2014.
- [29] B. Liang, W. You, L. Liu, W. Shi, and M. Heiderich, "Scriptless Timing Attacks on Web Browser Privacy," *Proc. of IEEE/IFIP DSN*, 2014.
- [30] N. Gelernter and A. Herzberg, "Cross-Site Search Attacks," *Proc. of ACM CCS*, 2015.
- [31] Y. Jia, X. Dong, Z. Liang, and P. Saxena, "I Know Where You've Been: Geo-Inference Attacks via the Browser Cache," *IEEE Internet*

Comput., vol.19, no.1, pp.44–53, 1 2015.

- [32] C. Jackson, A. Bortz, D. Boneh, and J.C. Mitchell, “Protecting Browser State From Web Privacy Attacks,” *Proc. of WWW*, 2006.
- [33] J. Ruderman, “css on a:visited can load an image and/or reveal if visitor been to a site,” https://bugzilla.mozilla.org/show_bug.cgi?id=57351, 2000.
- [34] D. Baron, “Preventing attacks on a user’s history through CSS :visited selectors,” <https://dbaron.org/mozilla/visited-privacy>, 2010.
- [35] W3C, “5.11.2 The link pseudo-classes: :link and :visited,” <https://www.w3.org/TR/CSS2/selector.html#link-pseudo-classes>, 2011.
- [36] P. Eckersley, “How Unique Is Your Web Browser?” <https://panopticklick.eff.org/static/browser-uniqueness.pdf>, 2009.
- [37] F. Roesner, T. Kohno, and D. Wetherall, “Detecting and Defending Against Third-Party Tracking on the Web,” *Proc. of USENIX NSDI*, 2012.
- [38] G. Acar, C. Eubank, S. Englehardt, M. Juarez, A. Narayanan, and C. Diaz, “The Web Never Forgets: Persistent Tracking Mechanisms in the Wild,” *Proc. of ACM CCS*, 2014.



Takuya Watanabe received M.E. degree in computer science and engineering from Waseda University, Japan in 2016. Since joining Nippon Telegraph and Telephone Corporation (NTT) in 2016, he has been engaged in research of consumer security and privacy. He is now with the Cyber Security Project of NTT Secure Platform Laboratories.



Eitaro Shioji received his B.E. in computer science and M.E. in communications and integrated systems from Tokyo Institute of Technology in 2008 and 2010, respectively. Since joining NTT in 2010, he has been engaged in R&D on computer security. His research interests include software security, program analysis, and reverse engineering.



Mitsuaki Akiyama received his M.E. and Ph.D. degrees in information science from Nara Institute of Science and Technology, Japan in 2007 and 2013. Since joining Nippon Telegraph and Telephone Corporation (NTT) in 2007, he has been engaged in research and development on cybersecurity. He is currently a Senior Distinguished Researcher with the Cyber Security Project of NTT Secure Platform Laboratories. His research interests include cybersecurity measurement, offensive security, and usable

security and privacy.



Keito Sasaoka received B.E degree in computer science from Waseda University in 2017. He is currently a 2nd-year graduate student in the Department of Computer Science and Engineering, Waseda University. His research interest is web privacy.



Takeshi Yagi received his B.E. degree in electrical and electronic engineering and his M.E. degree in science and technology from Chiba University, Japan in 2000 and 2002. He also received his Ph.D. degree in information science and technology from Osaka University, Osaka, Japan in 2013. He joined the Nippon Telegraph and Telephone Corporation (NTT) in 2002 and transferred to NTT Security (Japan) KK in 2018, where he is currently researching honeypots, security-data analysis based on machine learning, and security intelligence technologies such as URL/domain/IP blacklisting and reputation. He is a member of the Institute of Electrical and Electronics Engineers (IEEE) and the Institute of Electrical Engineers of Japan (IEEJ) and IEICE.



Tatsuya Mori is currently an associate professor at Waseda University, Tokyo, Japan. He received B.E. and M.E. degrees in applied physics, and Ph.D. degree in information science from the Waseda University, in 1997, 1999 and 2005, respectively. He joined NTT corporation in 1999. Since then, he has been engaged in the research of measurement and analysis of networked systems and network security. From Mar 2007 to Mar 2008, he was a visiting researcher at the University of Wisconsin-

Madison.