BakingTimer: Privacy Analysis of Server-Side Request Processing Time

Iskander Sanchez-Rola
University of Deusto
Symantec Research Labs

Davide Balzarotti
EURECOM

Igor Santos
University of Deusto

ABSTRACT

Cookies were originally introduced as a way to provide state awareness to websites, and are now one of the backbones of the current web. However, their use is not limited to store the login information or to save the current state of user browsing. In several cases, third-party cookies are deliberately used for web tracking, user analytics, and for online advertisement, with the subsequent privacy loss for the end users.

However, cookies are not the only technique capable of retrieving the users' browsing history. In fact, *history sniffing* techniques are capable of tracking the users' browsing history without relying on any specific code in a third-party website, but only on code executed within the visited site. Many sniffing techniques have been proposed to date, but they usually have several limitations and they are not able to differentiate between multiple possible states within the target application.

In this paper we propose BakingTimer, a new history sniffing technique based on timing the execution of server-side request processing code. This method is capable of retrieving partial or complete user browsing history, it does not require any permission, and it can be performed through both first and third-party scripts. We studied the impact of our timing side-channel attack to detect prior visits to websites, and discovered that it was capable of detecting the users state in more than half of the 10K websites analyzed, which is the largest test performed to date to test this type of techniques. We additionally performed a manual analysis to check the capabilities of the attack to differentiate between three states: never accessed, accessed and logged in. Moreover, we performed a set of stability tests, to verify that our time measurements are robust with respect to changes both in the network RTT and in the servers workload.

CCS CONCEPTS

• **Security and privacy → Browser security**.

KEYWORDS

user privacy; browser cookies; history sniffing

ACSAC '19, December 9–13, 2019, San Juan, PR, USA

ACM ISBN 978-1-4503-7628-0/19/12…$15.00
https://doi.org/10.1145/3359789.3359803

ACM Reference Format:
Iskander Sanchez-Rola, Davide Balzarotti, and Igor Santos. 2019. BakingTimer: Privacy Analysis of Server-Side Request Processing Time. In *2019 Annual Computer Security Applications Conference (ACSAC '19), December 9–13, 2019, San Juan, PR, USA.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3359789.3359803

1 INTRODUCTION

The World Wide Web (WWW) is built on top of the HyperText Transfer Protocol (HTTP), a stateless request/response protocol in which each request is executed independently from any other received before. However, most web applications need to keep track of the user's progress from one page to another. For instance, after a user has successfully logged into a service, she expects the application to remember her authentication and allow her to perform other actions in the same website accordingly.

In order to solve this problem, in 1994 Netscape Mosaic introduced the use of Cookies [42]. HTTP cookies are small fragments of data that servers can send to the users inside their responses (by using the `Set-Cookie` header field) or by the use of JavaScript code executed in a webpage in order create a cookie on the client-side (by invoking the `document.cookie` function). Either way, the browser stores the value of each cookie and includes them in every future request made to the same server. Today, cookies are used for a variety of different purposes, including to maintain the users' login status, to store different options that a user makes while browsing a website (such as the language preference or the acceptance/rejection of a specific consent), or to simply keep track of previous visits from the same user.

The cookies we described so far are called *first-party* cookies, as they are created by the website the user is visiting. However, these are not the only cookies that may be created in the browser. Websites load different types of resources to offer their services and the requests made to retrieve these resources may also trigger the creation of cookies. In many cases, these *third-party* cookies are used to track users among the different websites they visit. For instance, several studies have measured that a big percentage of websites on the Internet perform some form of user tracking [1, 13, 33, 35].

While tracking based on third-party cookies is one of the main techniques different companies use to offer personalized advertisements, other alternatives exist — for instance based on fingerprinting the browser or the user's machine by collecting either hardware or software information [7, 27, 30, 36]. These approaches can completely bypass all the browser protections regarding basic tracking, but their fingerprinting code needs to be executed in all the websites the user visits. Therefore, for websites that are not part of the main tracking networks, there is another option based on the so-called

history sniffing. History sniffing attacks can track the user without relying on any code executed in other third-party websites, just the one executed in the website accessed by the user. While many methods have been proposed in this line of research [6, 16, 26, 29, 45], most of them suffer from several severe restrictions. For instance, previous approaches only provide a coarse-grained classification (e.g., logged-in vs not logged-in), can be easily defeated by users without serious side-effects (e.g., by deleting the browser cache), and were all tested only in a handful of anecdotal cases. As an example, the two most recent works in this area, published in 2015 at the ACM CCS [45] and NDSS [29] conferences, were only evaluated on, respectively, five and eight target websites.

In this paper, we present a new timing side-channel attack, called BakingTimer, that only relies on the presence of first party cookies set by the target websites (which are therefore considered third-party in the context of the attacker page). Our system is based on the analysis of the time spent by the server to process a HTTP request, and by using this information is able to detect both if the user previously visited the website and whether she is currently logged in into it.

We then performed a set of experiments to measure how many websites are vulnerable to this attack. First, we checked if our methods was able to detect website accesses, and found that our prototype was able to detect the access in to more than half of the websites we analyzed. We tested our solution on over 10K websites belonging to different categories, resulting in the largest evaluation of a history sniffing technique performed to date. Second, we follow a similar approach as previous work, and manually tested the login detection capabilities in a small set of websites.

2 BACKGROUND

To ease the comprehension of our contribution, in this section we review several aspects regarding current cookie management, as well as different attacks presented in the literature. We then discuss the threat model we consider for our history sniffing attacks.

2.1 Browser Cookies

Cookies were introduced by Lou Montulli [38] while working in Netscape, and are considered the first method to track users on the Web. Cookies allow services to remember a particular user by storing snippets of data on the users' computers, maintaining a concrete browsing state for a returning visitor. After cookies were first presented, they were rapidly embraced by the web community because of their flexibility and the increased usability they enabled in a broad range of websites. As a result, they are now playing a core role as part of the Internet technology [33].

While cookies were not specifically designed to track users across websites, the abuse of their stateful nature started shortly after their first appearance. Since websites are composed of different resources that may be stored in the same domain hosting the page as well as in other third-party servers, these external resource providers have the capability of including cookies along with their provided resource. Therefore, if a third-party server provides resources to a large number of websites, it can certainly gather an accurate picture of a user's browsing history and profile her habits. For example, a website `site.com` includes an image from a third-party domain called `advertisement.com`. The server `advertisement.com` sends the image along with a `HTTP Set-Cookie` header, that will be stored on her machine. When the same user visits another website `site-two.com` that also utilizes images from `advertisement.com`, her browser will send the previously set cookie to `advertisement.com` alongside the request for the image. This allows the advertisement server to recognize the user and collect a list of the websites she regularly visits. This behavior, called *third-party cookies*, is arguably the most widespread technique for web tracking.

However, even the most privacy-invasive third-party cookies are an important part of the web. In fact, the most common usage of this specific type of cookies is web advertisement or analytics. Their web tracking capability is used to track users' browsing history and use this information to build a user profile for custom advertisements. In fact, advertising companies have already stated that web tracking is required for the web economy [40]. However, other techniques exist that can be used for analytics and targeting while preserving users' privacy (e.g., [2–4, 19, 22, 44]).

However, both third-party and first-party cookies are widely used and browsers even encourage their usage in order to ease usability. For example, when a user of the well-known Chrome browser decides to erase all the stored cookies, the browser warns the user that "*This will sign out of most websites*". Even more important are the settings regarding cookies. Chrome recommends to enable the permission for websites to read and write cookie data, and permits to block third-party cookies. Therefore, the importance of cookies is not only acknowledged by websites, but also by the browsers themselves albeit the privacy issues that may arise.

Even though it is not a common action among the average web users, *third-party cookies* can be blocked (as detailed in Section 2.1). However, browsers do not allow users to remove just these cookies, leaving as only option to manually remove them one by one.

2.2 History Sniffing

A large variety of history sniffing methods exists. We can group these techniques in two categories: CSS-based and timing-based.

CSS-based attacks. A common trick a page can use to detect if a user has visited other websites (out of a predefined list of targets) is to check the CSS style used to render links. For instance, by checking the `CSS:visited` style of a particular link [11] it is possible to tell if that link had been visited before (typically as it is displayed in a different color). Similarly, it is possible to craft other CSS-based techniques [23, 24, 41, 47] even using filters to exploit the differences of the DOM trees [46] by using a timing side-channel attack.

Timing-based attacks. There is a broad range of history-stealing techniques based on timing information. These approaches were first introduced for this purpose in order to compromise users' private data by measuring the time differences in accessing different third-party content [16], by discovering if it had been cached by the browser. Extracting users' true geolocalization for the same purpose is also possible through web timing, due to the high customization present in current websites. It is also possible to detect if the user

```php
<?php
$userID = "0bc63ecec05112d03544fde0b5a18c70";

if (isset($_COOKIE[["consent"]) {

    if (isset($_COOKIE[["userID"]) {

        if ($_POST["userID"] == $userID) {
            getUserData(); // Case C
        }

    } else {
        saveNavigation(); // Case B
    }

} else {
    askConsent(); // Case A
}
?>
```

Figure 1: Example code of a PHP server presenting the three different possible cases of a cookie process schema.

is currently logged in into certain websites by timing of specific requests [6], exploiting the AppCache [29], or by estimating the size of certain resources [45]. As explained in Section 3, our own attack belongs to the this timing-based attack category.

Countermeasures and Shortcomings Some of the attacks discussed above are already mitigated by browser vendors. For instance, for the CSS:visited style check, all the corresponding browser functions (e.g., getComputedStyle and querySelector) have been modified to always return that the user has never visited the link [31]. Despite these mitigations, recent work has shown that the attack is still possible using new features available in modern browsers. However, several possible defenses exist to avoid the problem, such as the ones proposed by Smith et al. [41]. In fact, one of these new techniques has already been blocked in recent versions of Google Chrome [20].

In fact, all existing techniques fall in the classic "arms race" category, in which attacker and researchers constantly discover new tricks that are in turn mitigated by browser vendors, website developers, or even simply careful user settings. Therefore, we decided to investigate if it was possible to devise a new technique that would 1) rely only on server-side information, and 2) that could not be easily prevented without degrading the performance or functionalities of a web application.

2.3 Threat Model

In the timing attack presented in this paper, we adopt the same threat model used by previous work in the area [6, 45]. In particular, we assume an attacker can run JavaScript code on the client browser to perform cross-origin requests. This code can be either loaded directly by the first-party website, or by a third-party service (e.g., by an advertisement or analytics company).

The information collected by our technique allows an attacker to determine which websites were previously visited by the user and on which website the user is currently logged in. There are multiple usages for this data that can result on serious security and privacy implications. The most obvious is related to advertisement, as the usage of the browsing history allows to display targeted advertisements. Moreover, an interested tracker could create a predefined list of websites and generate a temporal fingerprint of various users, indicating the user's state in each of them. Even if the fingerprint could not be used as an standalone fingerprinting solution, it will definitely improve the fingerprinting capabilities of other web tracking techniques. Finally, from a security point of view, this information can be used to perform targeted attacks against particular victims.

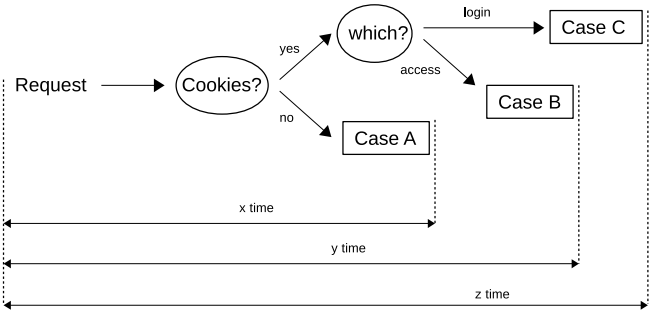

Figure 2: Server cookie procress schema.

3 BAKINGTIMER

Out of all the requests a web server receives, some contains cookies and some do not. The main observation behind our approach is that, when the browser sends a cookie along with an HTTP request, it is reasonable to assume that the server-side application will check its value, maybe use it to retrieve the associated user session and load additional data from the database, or that it will simply execute a different execution path with respect to a request that does not contain any cookie (in which case, for example, the application may execute the routines necessary to create new cookies).

For instance, Figure 1 shows a simple PHP skeleton that empathizes three different possible cases. First, the program checks if the user already accessed the website before, by testing for the presence of the consent cookie using the isset function. If the cookie is not found (either because it is the first time the user access the website or because it has been deleted), the program takes the path we named *Case A* that calls the askConsent function. Otherwise, the program performs additional checks by looking for some login information stored in the cookies. If userID is not indicated, *Case B* is followed, that calls function saveNavigation, else, the userID information is first validated, and then the application follows *Case C* by calling the getUserData function.

Figure 2 shows a simplified representation of the control-flow of our toy application, emphasizing the different paths. Our hypothesis is that the time difference among the three cases is sufficient, even across a network connection, to tell the three behaviors apart and therefore to accurately identify whether or not cookies are submitted alongside an HTTP request. In particular, there are two main tests that make detecting a timing difference possible: the first to verify if there are cookies at all and the second to analyze them and load session data. While the comparison themselves are too

fast to be identified, the different functions invoked in the three cases may not be. In our toy scenario, these results in the request take x seconds to be processed if no cookies are found, y seconds if they exist but there is not a current active session, and z seconds if the user is currently logged in.

Our prototype tool, called BakingTimer, performs cross-origin requests towards a number of target websites. Each request is performed multiple times, with and without cookies. In fact, while JavaScript code cannot directly test for the existence of cookies belonging to other domains, it can issue HTTP requests in two different ways: (i) by letting the browser send cookies (if any exist for the target domain) or (ii) by forcing the browser to make a request without cookies. So, in one case the code knows that no cookie was present in the request, while in the other cookies may or may not be present depending whether they previously existing in the user browser. The time difference between the two requests, issued multiple times to account for small noise effects, is then used by our tool to detect if the browser already had cookies for that particular target. Even if some cookies can be created when performing these requests, our experiments show that in most cases either the number or type of these cookies differ from what it is set when the user visits the website. We will discuss this aspect in more details in Section 4 and Section 7.

3.1 Retrieval Phase

The first phase of our approach is the retrieval phase, in which the algorithm computes the time required by a server to execute different paths, according to hypothetical cookies sent by the browser. A simplified pseudo-code of the algorithm is presented in Figure 3. Note that the full implementation is executed performing asynchronous/non-blocking requests and using `onreadystatechange` event handlers. Moreover, even if security policies like *Same-Origin Policy* (SOP) or *Cross-Origin Resource Sharing* (CORS) may limit the capabilities of interacting with the response of cross-origin requests, the event handler would still be able to correctly catch when the response arrives to the browser – making our attack possible in these cases.

To measure the round-trip times (RTTs), we use the *User Timing API* [12] implemented in every major browser. The best solution to get a measure of time independent of the current CPU load of the server or the current speed and congestion of the network is to perform a comparison with two different requests sent at the same time. Both time information are obtained side to side, so the workload of the network will likely be roughly the same in the two cases (especially when the experiment is repeated multiple times). To make it more clear, let us consider a simple example. Imagine we are timing the execution of certain function of the server that directly depends on the existence of a specific cookie. Our system would execute a measurement by executing in parallel two requests, with and without the cookies. In a first measurement, when the network workload is low, the system can obtain a value of 30.5ms when cookies are sent and a value of 20.3ms without cookies — thus estimating the execution time of the function in 10.2ms. In a second repetition of the test, when maybe the network load is higher than before, the same experiment can return two different values, such as 45.1ms with cookies and 34.7ms without — leading to

Input: n the number of comparisons to perform.
Output: cs an array of arrays of numbers representing the server request processing schema: each position are the result of timings with and without cookies.

```
1  Function bakingTimer (n)
2      i ←1;
3      cs ← float[][] of size n×2;
4      while i ≤ n do
5          j ← 1;
6          while j ≤ 2 do
7              startTime ← GetCurrentTime() ;
8              if j % 2 = 0 then
9                  Request(cookies);
10             else
11                 Request();
12             end
13             endTime ← GetCurrentTime();
14             logTime ← endTime − startTime;
15             cs[j][i] ← logTime;
16             j ← j + 1;
17         end
18         i ← i + 1;
19     end
20     return cs;
```

Figure 3: BakingTimer simplified retrieval pseudo-code (implemented using asynchronous/non-blocking requests and `onreadystatechange` event handlers).

a similar (even if not exactly the same) time estimation. In Section 5 we present a number of *stability tests*, designed to verify to which extent our technique is robust against external factors that could tamper the results.

Following the schema presented in Figure 2, if a request requires x *seconds* to be processed (as opposed to y or z seconds), a simple comparison can be used to calculate the actual state of the user in relation with the website. The first request, the one we named as *Case A*, is issued without cookies by using:

```
xmlHttpRequest.withCredentials = false;
```

This guarantees that the server would not receive cookies, even if the browser would have normally sent some for that particular website. Then, our code repeat the request, this time by setting

```
xmlHttpRequest.withCredentials = true;.
```

It is important to remark that the cookies sent in this second request (if any) are completely invisible for our Javascript code. However, the code can still measure the time required to receive an answer, and use this information to infer the previous relation between the user and the target website: if the cookies were created in a simple access, the server response time will fall in *Case B*, but if some login cookies are stored in the browser the time would be more consistent with *Case C*.

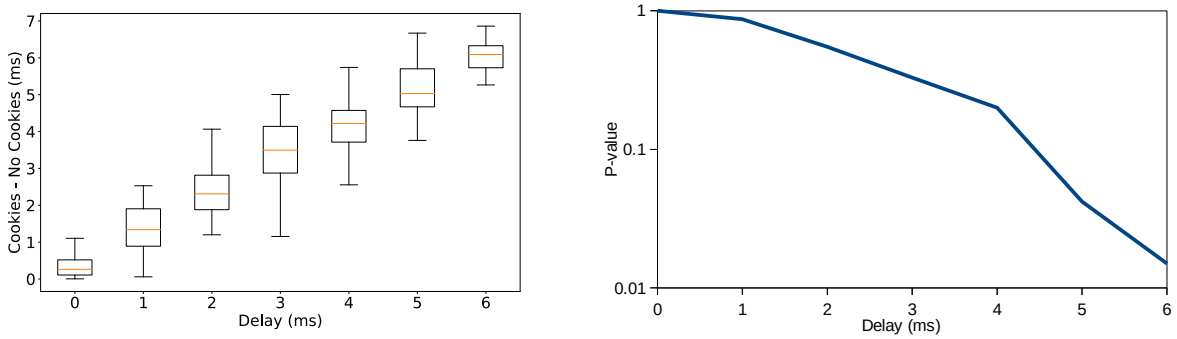

(a) Time deltas at differentdelays introduced in the server. **(b) P-value (T-test) at different delays introduced in the server.**

Figure 4: BakingTimer resolution test.

To summarize:

- **Not Accessed:** If the user never accessed the website, and we perform the timings described above, we will obtain a result close to zero. In fact, the first request will take x seconds, because we made the request without cookies. However, also the second request with cookies will take roughly the same amount of time, as no cookie were available in the browser for the target. In this situation we can infer than both calls were the same, indicating no previous access to the website under inspection.
- **Accessed/Logged:** If the user accessed the website in the past or if it is currently logged in into it, the result of the comparison of the time taken by the two request will be different from zero. As we are not using absolute values, but only time differences between consecutive requests, our technique can compensate for differences due to different network delays or workloads.
  A more fine-grained classification to distinguish between previous access and current login is also possible if the attacker has information about the different time required for specific actions. We will investigate this option and provide more details in the comparison phase section.

The algorithm (see Figure 3) takes one parameter n that indicates the number of comparisons to perform. As servers can have specific changes in their workload for many different reasons independent of our tests, we decided to obtain more than one simple comparison. Clearly, the more comparisons the attacker is able (or willing) to perform, the more precise is the model of the server-side computation she can obtain. Nevertheless, there is an obvious trade-off between the time spent in fingerprinting a singe website and the number of websites the attacker wants to test. We will investigate different options and the impact of the number of comparisons on the accuracy of the classification in Section 4.

3.2 Comparison Phase

As explained in the previous section, our code performs a number of measurements by timings different HTTP requests for a target website. By subtracting each pair of requests, with and without cookie, we obtain an array of n time deltas.

In order to classify this data, we need to obtain a ground truth dataset for the selected websites. To this end, we need to retrieve the computing time information for each of the cases we are interested to detect. By using this information, we can then statistically test whether a set of measurements belong to one group or another, simply by computing a T-test over the two data samples. A T-test is a two-sided test for the null hypothesis that two samples have identical average values. The result tells us which of the three presented cases does a certain experiment belongs to.

3.3 BakingTimer's Resolution

Before we started our large scale experiments, we wanted to experimentally verify that our hypothesis is correct and that the time difference among different server functionalities can be successfully distinguished over a long-distance network connection. For this reason we implemented a toy service based on `web.py` [43], a simple but powerful Python framework that has been used by famous websites such as Reddit and Yandex. In our example, we controlled the time spent in each path. All HTTP requests sent to the service were issued from a browser located in a different country of where the server was located (Spain and France respectively), to avoid any possible bias introduced by short-distance connections. The average ping time among the two machines was 13.6 milliseconds.

The server-side application consist in less than 10 lines of code with just a single Python function. In this function, we receive the GET requests, and using the `web.cookies` method, we are able to check if the request includes any cookie, and its value. The service was designed to introduce no delay for requests without cookies, and a configurable delay (ranging from 1 to 6 milliseconds in our tests) to the processing of any request containing cookies. The specific delay was indicated using that same cookie, which was created with JavaScript thought `document.cookie`. We were able to control the introduced delay invoking the function `sleep` from the Python `time` library.

The results of our experiments are summarized in Figure 4. The graphs show that it is possible to detect the time the server spent

processing each request quite precisely (see Figure 4a), despite the network load. However, with delays below four or five milliseconds, the difference among the two sets of requests measured by the browser is not statistically significant and therefore an attacker could not conclude whether or not cookies were present in the browser for the target website (see Figure 4b). Instead, if the difference between two paths was equal or above five milliseconds, then even over a long distance connection, it was possible to reliably tell the difference between the two cases. This indicates that it is not necessary for a website to perform complex computations to be vulnerable to our technique and even looking up some data from a database may be sufficient to make it vulnerable to our attack. Obviously, less optimized servers are more prone to incur into larger delays that are easier to observe remotely, while high-end servers may make the path detection more difficult and error prone. We will analyze all these situations in the following section.

4 EXPERIMENTS

In most of the previous studies on history sniffing, the authors limited their experiments to just few selected websites. This poor coverage is mainly due to the fact that existing solutions could generally only distinguish between two states: currently logged in or not logged in. As a result, experiments required the authors to manually create accounts for different services, thus reducing the number of websites that could be tested.

Our technique allows instead to distinguish among multiple states, and in particular to differentiate between websites that have been previously visited by the victim from those that have not. Therefore, on top of performing a login detection case study (detailed in Section 6), we also conducted a large scale experiment to measure the percentage of different websites that are vulnerable to our attack.

The dataset used in our tests consists of two different groups: (i) highly accessed websites and (ii) websites related to sensitive information that users may want to keep private. For the first group we selected websites from the Alexa [39] Top5K list, to verify whether extremely popular services are also vulnerable to our timing side-channel attack. The second group is composed instead by websites that can be used to detect private personal information of the user, such as medical or religious websites. We used categories defined as sensitive in various data protection laws [14, 15, 28]. Since many of the entries in this group are not included in the Alexa Top1M, we could also check if not highly accessed websites are more or less vulnerable than highly accessed ones.

4.1 Domain Selection

We first populated our list of personal information websites by performing various queries obtained through the auto-complete option offered by different search engines (e.g., *"cancer treatment side effects"*). We made five different queries for the following six categories: medical, legal, financial, sexual identity, political, and religion. Our tool issues each query on four search engines (Google, Bing, Yandex, and DuckDuckGO) and retrieved the top 100 results from each of them.

To avoid an overlapping between the two lists, we removed from this group domains that also belonged to the Alexa Top10K list. We also removed any website that appeared in multiple categories, as we wanted to focus only on specific areas in isolation. Finally, we obtained a set of 5,243 unique personal information websites. In order to balance the two groups in the dataset, we selected the same number of websites from the Alexa top list. The combination of the two groups resulted in a final dataset of 10,486 websites.

4.2 Methodology

We implemented our BakingTimer proof-of-concept by using a custom crawler based on the well-known web browser Chrome, without using any flag that may influence the network connectivity in any way. We can perform multiple actions using the Chrome debugging protocol (with the `remote-debugging-port` parameter), which allows developers to control the browser [8]. For instance, we can access different websites using the `navigate` command, or run JavaScript code on the current website once loaded with the `loadEventFired` event and the `evaluate` command. The executed code uses `xmlHttpRequest` to perform the corresponding requests of the presented technique. Even if it would have been possible to implement the experiments by using a simple Python script, we decided to rely on a real-world browser to obtain the most realistic results possible. Our crawler follows two different phases in order to obtain the data that we will later used to check if a target server is actually vulnerable to the attack.

- **Never Visited:** First, the tool cleans every previous access information stored in the browser. Then, it starts making the different requests described in Section 3 (i.e., both with and without cookies, which in this case are none) from a third-party website (blank tab). This data will be later used as a baseline for requests performed when no previous access was done.
- **Previously Visited** In this case, the browser first accesses the website under inspection. No cleaning process is performed, so all cookies automatically created by the website are saved in the browser and sent to the server in the following requests. After that, it goes to a third-party website (blank tab) and starts making the different requests described in Section 3 to that same website under inspection.

Once all the data was retrieved, we performed the statistical tests described in Section 3 in order to identify whether the time information in the two groups of requests are statistically different or not. We also repeated the experiment with different number of requests in order to check their influence on the final result. To be practical, we tested with a minimum of 10 and a maximum of 50 comparisons (therefore ranging from 20 to 100 HTTP request per target website). The higher the number the more stable is the measurement, but so is the amount of time required to run the test. Therefore, the choice boils down to a trade-off between precision and scalability. We believe that if the attacker is only interested in a handful of websites, then it would be possible to perform even more than 50 comparisons.

It is also important to note that since we need to repeat the requests multiple times, it is possible that the first request "pollutes" the browser by setting cookies that are later sent in the following tests. In other words, if i) the target website sets cookies on cross-origin requests, and ii) those cookies are the same (in number and

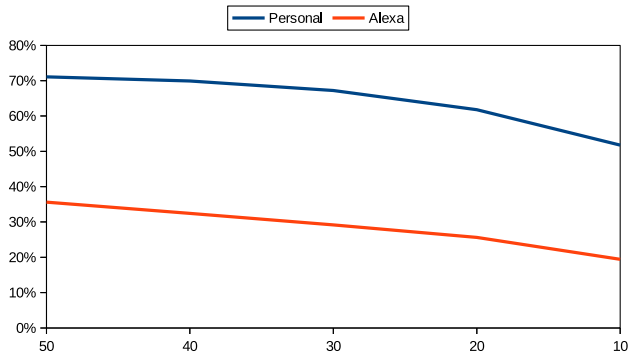

Figure 5: Percentage of vulnerable websites depending in the number of comparisons performed.

nature) of those set when the website is loaded in the browser, then our timing attack would not work.

However, this does not seem to be very common, as more than half of the 10,486 websites we tested are vulnerable to our technique. The actual percentage varies between 40%, when the minimum number of comparisons is performed, and 53.34% if our attack performs 50 comparisons. Figure 5 shows the success rate at different numbers of comparison for the two groups separately.

4.3 Highly Popular Websites

This group includes 5,243 websites from the Alexa top websites list. As websites in these categories are visited by a large number of users, most of their servers and the code they run are likely more optimized, thus making more difficult to measure the timing side channel. This is confirmed by the fact that our attack worked on 35.61% of the websites in this category. This result is still quite severe, as it means than a webpage could still reliably identify if its users had previously visited more than one third of the most popular pages on the Internet.

In order to get a deeper analysis of the obtained results, we clustered the websites in different categories and we computed the percentage of each of them that were vulnerable to our attacks. To determine the category, we used three services: Cloudacl [10], Blocksi [5], and Fortiguard [17]. Their category names are similar, and, after a normalization process, we settled for 78 different category names. Table 1 shows that the top 6 categories vulnerable to the attacks include around 40% of the websites, with a peak of 43.75% in the case of sport-related websites in the Alexa top list.

4.4 Privacy-Sensitive Websites

This group includes 5,243 websites from six different categories related to private personal information — i.e., medical, legal, financial, sexual identity, political, and religion. The results from these websites allows us to understand two different aspects. First, we can verify the amount of websites directly related to sensitive information that are vulnerable to our attack. Second, it gives us an opportunity to test less popular websites (as 85% of the vulnerable sites in this category are ranked below the Alexa Top500K) to observe whether smaller server infrastructures can result in a higher accuracy of our technique.

Table 1: Percentage of websites vulnerable to our attack (top six private-sensitive on the top half and top six highly popular in the bottom).

Category	% Vulnerable
Medical	72.63
Religion	71.66
Financial	71.63
Political	70.73
Sexual Identity	70.38
Legal	69.39
Sports	43.75
Search/Portal	42.25
Government	40.63
Travel	40.40
Gaming	39.39
Adult/Pornography	39.06

The results of our experiments show that a stunning 71.07% of all the analyzed websites in this group are vulnerable to the BakingTimer attack. If we break down the result by category, we see that all have similar percentages and there is no clear difference between them (see Table 1). This result is much higher than the one obtained in the top Alexa group, but the difference is not due to the number of cookies. In fact, we compared the mean and standard deviation of the number of cookies in privacy-sensitive websites and highly popular websites. Unsurprisingly, the results show that highly accessed websites have a higher mean number of cookies (9.03±7.87) compared to the number of cookies in private personal information websites (5.83±5.49). This means that the main reason behind the difference is not linked to the number of cookies, but more likely to the slower servers or the less optimized code responsible to process the incoming requests.

5 STABILITY TEST

Our attack relies on a time-based side channel used to distinguish among different execution paths on the server-side application code. The fact that BakingTimer does not look at absolute times but at the difference between two consecutive requests minimizes the effects of network delays on our computation. However, even if the time difference between the two request is minimal, it is possible that small non-deterministic variations such as the jitter, bandwidth, or network congestion and routing can introduce a certain level of noise in the measurement. To account for these small fluctuations, BakingTimer needs to repeat each test multiple times. As shown in Section 4, ten comparisons are sufficient to detect nearly 40% of all the analyzed websites, which increases to over 50% if we perform 50 comparisons.

In order to obtain a clear view of the specific effect the network or the server load can have in our measurements, we decided to perform a dedicated stability test. For this experiment we randomly

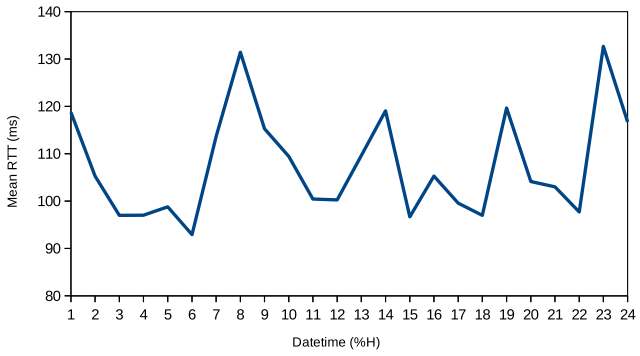

Figure 6: Mean RTT of one website never visited before, during a full day (with data every hour).

picked 25 website detected as not vulnerable and 25 websites detected as vulnerable to our attack. Following the same approach presented in the general experiment, we now repeated our test every hour for a period of 24h for each website. Moreover, to be more realistic, we computed the ground truth on one day, and performed the attacks on the following day. This resulted in a total of 48 checks per websites, and 2,400 tests globally.

From the group of the 25 websites not vulnerable to our attack, all our tests returned the same result (which confirmed the absence of the side channel). This results proves the stability of the presented method from a network perspective. More concretely, regarding fluctuations, Figure 6 shows, for one of the websites analyzed, the different mean RTTs we registered each hour. Even if there were considerable fluctuations on the network speed during the day, we were still able to perform the correct classification.

From the group of the 25 vulnerable websites, we were able to correctly identify when each website was previously visited – in all our tests. Instead, in the case in which we did previously visited the websites, there was one case in which (at 21 and 22pm), we incorrectly identified a single website as visited. Nevertheless, in total from the 2,400 checks performed in this experiment, we only incurred in two false positives and no false negatives, indicating the high reliability of the presented history sniffing technique.

6 LOGIN DETECTION

In this section we look at performing a more fine-grained classification of the user history, not just by telling if a victim has visited a target website in the past, but also to distinguish a simple visit from owning an account or being currently logged in into the service.

Since this analysis requires a considerable manual effort to set up the accounts and operate the websites, we will limit our study to few examples taken from the two groups in our dataset. This approach is similar to the one regularly adopted by other related work on history stealing techniques [29, 45]. It is important to remark that we did no test websites that required a registration fee or that had a complex registration procedures (e.g., requiring a phone number verification or a bank account).

this cases, a third-party website can check if the user ever accessed any of the websites under attack, and can even check if the user is logged in. This type of information could be extremely useful for a malicious attacker. For instance, it could be used to perform targeted phishing attacks against users, to steal the login credentials of those affected websites. Moreover, it will also be beneficial for other types of attacks like Cross-Site Request Forgery (CSRF), in which the attacker could use the state of the user in a website vulnerable to this type of attacks to perform privileged actions on the users accounts, such as password changes or information retrieval [9, 34]. Actually, the attacker does not need to control the third-party website, just have his code executed on them, for example via an advertisement (as explained in Section 2).

6.1 Highly Accessed Websites

From this category we selected two popular websites, related to gaming and clothing, that were detected to be vulnerable to our BakingTimer attack (Section 4), and that have been the target of different phishing attacks — in one case for the huge number of users and in the other case for the high economic value of their customers. More concretely, World of Warcraft (WoW) and Gucci.

In both cases, after the user logs in, a number of cookies are created in the user's browser to store this new status (e.g., `wow-session` and `gucci-cart` respectively). The presence of these cookies make the server take a different execution path, resulting in a different computation time. We followed the same analysis principles as the ones used when we analyzed the websites in Section 4, and detected both websites fall in our simplified three-paths model presented in Figure 2. The results show than each of the different states (e.g., not accessed, accessed, and logged in), does not match any of the other two states when performing the statistical test of the comparison phase (see Section 3).

6.2 Private Personal Information Websites

Detecting if a user has previously visited a website linked to specific information such as religion or sexual identity can leak some private information to third-parties the user may not even be aware of. However, if the third party could actually detect that the user is currently logged in into one of those websites, the link between the user and the website becomes much stronger.

From this category we picked a religious website (Dynamic Catholic) and a sexual related chat/forum (LGBTchat.net). Again, the presence of login cookies (i.e., `xf_user` and `frontend_cid`) made the two applications take different execution paths, whose computation time was sufficiently different to be reliably fingerprinted by our solution.

6.3 Persistent Login Information

In all previous cases, when the user logs out from the website, the different cookies related to the login status are deleted by the server. In this situation, a third-party website would still be able to detect that the user has visited the website in the past, but it would not be able to distinguish if she had an account and she ever logged into the service.

While this may seem obvious, there are also websites for which it is not true. In fact, some sites do not properly (or at least not

Table 2: A comparison of current state-of-the-art timing methods for history sniffing.

Approach	Type	Login Status	Difficult Clean	Previous Access	Websites Analyzed
Timing Attacks on Web Privacy [16]	Web & DNS Caching	✗	✗	✓	<10
Exposing private information... [6]	Server Boolean Values	✓	✓	✗	<10
Cross-origin Pixel Stealing [26]	Cross-origin DOM Structure	✓	✓	✗	<10
Identifying Cross-origin... [29]	HTML5 Application Cache	✓	✗	✗	<10
The Clock is Still Ticking [45]	Cross-origin Resource Size	✓	✓	✗	<10
Our Approach (BakingTimer)	Cookie-based Request Timing	✓	✓	✓	10,486

completely) delete all cookies they created in relation to the login process. This could be either because the cookie deleting process was not performed correctly, or because the website developers explicitly decided to maintain part of the login information stored in the cookies even when the user is not logged in. However, the presence of these cookies can be sufficient to trigger a different code execution in the application that can be detected by our technique. This allows third-party to be able to differentiate users that just accessed the websites from those who own an account, even if they are not logged in at the time the test is performed.

For instance both Microsoft/MSN and Openload fall into this category because of the presence of, respectively, a `MUID` cookie and a cookie with a `MD5` hash as name. In both cases, when the user logs out of the service, some cookies are deleted but some other are maintained in the browser. In this two specific cases, we first classified the websites following the same three-state schema as in previous cases. Then, we logged out of the service and checked if this state would be classified as logged or accessed. Our results show that in both cases, the comparison phases classified this state as logged. This indicates that even if the user logged out, if the websites does not correctly delete all related cookies, it would be possible to detect a previous logged state.

7 DISCUSSION

Even if a user trusts all the websites she visits, many websites include a large number of third-party services and resources to improve their usage or to monetize their traffic. All these scripts loaded on the website, including simple advertisement banners, can track users and perform many different attacks, including the one we presented in this paper.

7.1 Result Interpretation

In order to obtain the most reliable results, it is important to perform the experiments against multiple real-world websites. In fact, synthetic websites or small sample sets may not correctly capture all the implementation patterns encountered in the wild. Our tests show that more than half of the websites we analyzed are vulnerable to our attack. This means two important things. First, that there is a measurable, statistically significant difference between the time the server spend processing the two classes of requests (with, and without the access cookies). Second, that either the website does not set cookies on cross-origin requests, or that those cookies are different from the one created on a regular access.

7.2 Comparison with Similar Techniques

Table 2 shows the different state-of-the-art timing methods and their different characteristics, both in terms of adopted technique and on the type and size of the experiments performed to validate.

The majority of previous works allowed an attacker to detect the login status of a victim in other websites. Only one [16], apart form ours, allows also to detect the access state (but this same technique is unable to detect the login status). The only technique able to detect both access and login state is the one presented in this paper.

Most of existing attacks, including our own, do not rely on any other browser resource than the cookies. This makes the technique resilient to generic browsing history cleaning processes, as browsers explicitly discourages user to delete cookies in their settings (see Section 2). Two techniques are instead based on different types of browser caching that, on the contrary of cookies, are even deleted by default, and therefore users can easily and without any major consequence delete them when needed.

Regarding the number of websites analyzed, as nearly all the techniques are only able to detect the login status, the manual effort needed to perform a big scale analysis made large scale experiments unfeasible. We also presented a similar set of experiments in Section 6, but we also performed an automatic analysis of over 10k websites divided in different categories to provide a general overview of how effective this attack can be in the wild.

7.3 Countermeasures

There are two possible timing countermeasures that could be implemented to avoid different types of server-side attacks such as the one presented in this paper [32, 37]. One consists in including a random delay in the response time of the servers. However, this would just increase the noise and a larger number of comparisons may compensate for small random changes in the processing time. The other method would be to change the web application to have fixed response times for sensitive requests. On top of being very difficult to be properly implemented, this solution would also have a large performance impact — considerably reducing the number of requests per second a web site can sustain. Moreover, as the fixed time would need to be exactly the same for all the sensitive request, they should be as slow as the slowest response the server can make.

For all these reasons, we believe none of these mitigations are practical and feasible to implement in a real-world deployment. Like other time-based network fingerprinting solutions, BakingTimer is therefore very difficult to mitigate.

Another possible solution to the problem, would be related to the creation of the cookies themselves. Some browsers are going to start supporting the *SameSite* attribute [18, 21], thanks to which, websites can specifically indicate that they do not want a cookie to be send in third-party requests. This option is a very interesting approach, and can stop attacks similar to ours. However, in order to completely protect from the technique presented in this paper, all cookies must set this attribute. As long as one of the cookies involved does not indicate it, the attack would still work. On the other hand, these changes could impact the state detection of the attack differently. For instance, due to the sensitive nature of login cookies, they could be more prone to use this option. Nevertheless, its important to remark that some sites could lose part of their core functionalities as a result of using SameSite cookies. Many types of websites, such as social networks or cashback services, rely on cookies to be added in third-party requests, in consequence, its global applicability could be limited in some situations.

8 RELATED WORK

History sniffing attacks are a widely explored topic with different techniques and solutions presented over the years. Clover [11] found that it was possible to identify previously visited websites just checking the CSS:`visited` style of a specially crafted link though the `getComputedStyle` method in JavaScript. Many other similar attacks appeared using different CSS-based techniques [23, 24, 41, 47]. Kotcher et al. [26] discovered that besides from the above mentioned attacks, the usage of CSS filters allows the involuntary revelation of sensitive data, such as text tokens, exploiting time differences to render various DOM trees. Weinber et al. [46] followed another direction, using interactive techniques to get the information. While these attacks are much slower, the protection methods are in principle more difficult to implement.

With a different approach, and leaving CSS aside, Felten and Schneider [16] introduced web timing attacks as a tool to compromise users private data and, specifically, their web-browsing history. Particularly, they proposed a method based on leveraging the different forms of web browser cache to obtain user specific browsing information. By measuring the time needed to access certain data from a third-party website, the attacker could determine if that specific data was cached or not, indicating a previous access. Some years later, Jia et al. [25] analyzed the possibility of identifying the geo-location of a given visitor using to the customization of services performed by websites. As this location-sensitive content is also cached, it is possible to determine the location by checking this concrete data and without relying in any other technique.

Bortz et al. [6] organized JavaScript web timing attacks in two different types of attacks: (i) direct timing, based on measuring the difference in time of diverse HTTP requests and (ii) cross-site timing, that allows to retrieve private client-side data. The first type could expose data that may be used to prove the validity of specific user information in certain secure website, such as the username. The second attack type follows the same line of previous work by Felten and Schneider. They also performed some experiments that suggested that these timing vulnerabilities were more common than initially expected.

Two recent studies show that these attacks are far from being solved. Van Goethem et al. [45] proposed new timing techniques based on estimating the size of cross-origin resources. Since the measurement starts after the resources are downloaded, it does not suffer from unfavorable network conditions. The study also shows that these attacks could be used in various platforms, increasing the attack surface and the number of potential victims. The specific size of the resource can leak the current state of the user in the website. Lee et al. [29] demonstrated that using HTML5's AppCache functionality (to enable offline access), an attacker can correctly identify the status of a target URL. This information can later be used to check if a user is logged or not in certain website.

However, these timing techniques can generally only determine if the user is logged on a specific website or some isolated data, but not if she has just previously accessed it. Moreover, some of them use resources easily cleanable by the user, like different cache options, as they do not imply any visible consequence to the final user.

9 CONCLUSIONS

Many different threats against the users security and privacy can benefit from a list of websites previously accessed by the user and a list of services where the user is logged in or ever logged in.

In this paper, we show that simply using cookies of third-party websites, is possible the detect the specific state (e.g., accessed and logged) of a user in certain website, which outperforms previous techniques that are only able to detect one single state. In particular, we present a novel timing side-channel attack against server-side request processing schema. This technique is capable of detecting execution paths with more than 5 milliseconds of difference between each other.

We also analyzed real-world servers to detect the percentage of websites vulnerable to the presented attack. All previous work analyzed less than 10 websites (manually), as they generally only detect the logged status. We performed this same analysis, and additionally, we performed an automated check of 10k websites from different categories and number of users. Results show that more than half of the websites are vulnerable to our technique.

ACKNOWLEDGMENTS

This work is partially supported by the Basque Government under a pre-doctoral grant given to Iskander Sanchez-Rola.

REFERENCES

[1] Acar, G., Juarez, M., Nikiforakis, N., Diaz, C., Gürses, S., Piessens, F., and Preneel, B. FPDetective: dusting the web for fingerprinters. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)* (2013).

[2] Akkus, I. E., Chen, R., Hardt, M., Francis, P., and Gehrke, J. Non-tracking web analytics. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CSS)* (2012), ACM, pp. 687–698.

[3] Backes, M., Kate, A., Maffei, M., and Pecina, K. Obliviad: Provably secure and practical online behavioral advertising. In *Proceedings of the IEEE Symposium on Security and Privacy (Oakland)* (2012).

[4] Bilenko, M., Richardson, M., and Tsai, J. Targeted, not tracked: Client-side solutions for privacy-friendly behavioral advertising. In *Proceedings of the Privacy Enhancing Technologies (PETS)* (2011).

[5] Blocksi. Web content filtering. `http://www.blocksi.net/`, 2018.

[6] Bortz, A., and Boneh, D. Exposing private information by timing web applications. In *Proceedings of the 16th international conference on World Wide Web* (2007), ACM, pp. 621–628.

[7] Cao, Y., Li, S., and Wijmans, E. (Cross-)browser fingerprinting via os and hardware level features. In *Proceedings of the Network and Distributed System Symposium (NDSS)* (2017).

[8] ChromeDevTools. DevTools Protocol API. `https://github.com/ChromeDevTools/debugger-protocol-viewer`, 2019.

[9] Cisco Adaptive Security Appliance. CVE-2019-1713. `https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2019-1713`, 2019.

[10] Cloudacl. Web security service. `http://www.cloudacl.com/`, 2018.

[11] Clover, A. Css visited pages disclosure. *BUGTRAQ mailing list posting* (2002).

[12] Consortium, W. W. W. User timing. `https://www.w3.org/TR/user-timing/`, 2018.

[13] Englehardt, S., and Narayanan, A. Online tracking: A 1-million-site measurement and analysis. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)* (2016).

[14] Directive 2009/136/EC of the European Parliament and of the Council of 25 November 2009. *Official Journal of the European Union* (2009).

[15] Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *Official Journal of the European Union* (2016).

[16] Felten, E. W., and Schneider, M. A. Timing attacks on web privacy. In *Proceedings of the 2000 ACM SIGSAC conference on Computer & communications security* (2000), ACM, pp. 25–32.

[17] Fortinet. Fortiguard web filtering. `http://www.fortiguard.com/`, 2018.

[18] Foundation, M. Supporting same-site cookies in firefox 60. `https://blog.mozilla.org/security/2018/04/24/same-site-cookies-in-firefox-60/`, 2019.

[19] Fredrikson, M., and Livshits, B. Repriv: Re-imagining content personalization and in-browser privacy. In *Proceedings of the IEEE Symposium on Security and Privacy (Oakland)* (2011).

[20] Google. Leak of visited status of page in blink. `https://chromereleases.googleblog.com/2018/05/stable-channel-update-for-desktop_58.html`, 2018.

[21] Google. Improving privacy and security on the web. `https://blog.chromium.org/2019/05/improving-privacy-and-security-on-web.html`, 2019.

[22] Guha, S., Cheng, B., and Francis, P. Privad: practical privacy in online advertising. In *Proceedings of the USENIX conference on Networked Systems Design and Implementation (NDSI)* (2011).

[23] Heiderich, M., Niemietz, M., Schuster, F., Holz, T., and Schwenk, J. Scriptless attacks: stealing the pie without touching the sill. In *Proceedings of the 2012 ACM conference on Computer and communications security* (2012), ACM, pp. 760–771.

[24] Janc, A., and Olejnik, L. Web browser history detection as a real-world privacy threat. In *European Symposium on Research in Computer Security* (2010), Springer, pp. 215–231.

[25] Jia, Y., Dong, X., Liang, Z., and Saxena, P. I know where you've been: Geo-inference attacks via the browser cache. *IEEE Internet Computing 19*, 1 (2015), 44–53.

[26] Kotcher, R., Pei, Y., Jumde, P., and Jackson, C. Cross-origin pixel stealing: timing attacks using css filters. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security* (2013), ACM, pp. 1055–1062.

[27] Laperdrix, P., Rudametkin, W., and Baudry, B. Beauty and the beast: Diverting modern web browsers to build unique browser fingerprints. In *Proceedings of the IEEE Symposium on Security and Privacy (Oakland)* (2016).

[28] Lapowsky, I. California unanimously passes historic privacy bill. Wired, 06 2018.

[29] Lee, S., Kim, H., and Kim, J. Identifying cross-origin resource status using application cache. In *NDSS* (2015).

[30] Mowery, K., and Shacham, H. Pixel perfect: Fingerprinting canvas in HTML5. In *Proceedings of the Web 2.0 Workshop on Security and Privacy (W2SP)* (2012).

[31] Mozilla. Privacy and the :visited selector. `https://developer.mozilla.org/en-US/docs/Web/CSS/Privacy_and_the_:visited_selector`, 2018.

[32] Nagami, Y., Miyamoto, D., Hazeyama, H., and Kadobayashi, Y. An independent evaluation of web timing attack and its countermeasure. In *Availability, Reliability and Security, 2008. ARES 08. Third International Conference on* (2008), IEEE, pp. 1319–1324.

[33] Nikiforakis, N., Kapravelos, A., Joosen, W., Kruegel, C., Piessens, F., and Vigna, G. Cookieless monster: Exploring the ecosystem of web-based device fingerprinting. In *Proceedings of IEEE Symposium on Security and Privacy (Oakland)* (2013).

[34] phpMyAdmin. CVE-2019-12616. `http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2019-12616`, 2019.

[35] Sanchez-Rola, I., and Santos, I. Knockin' on trackers' door: Large-scale automatic analysis of web tracking. In *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA)*, pp. 281–302.

[36] Sanchez-Rola, I., Santos, I., and Balzarotti, D. Clock Around the Clock: Time-Based Device Fingerprinting. In *Proceedings of the ACM SIGSAC conference on Computer & communications security (CCS)* (2018).

[37] Schinzel, S. An efficient mitigation method for timing side channels on the web. In *2nd International Workshop on Constructive Side-Channel Analysis and Secure Design (COSADE)* (2011).

[38] Schwartz, J. Giving the web a memory cost its users privacy. `http://www.nytimes.com/2001/09/04/technology/04COOK.html`, 2001.

[39] Services, A. W. Alexa top sites. `https://aws.amazon.com/es/alexa-top-sites/`, 2018.

[40] Singer, N. Do not track? advertisers say "don't tread on us". `http://www.nytimes.com/2012/10/14/technology/do-not-track-movement-is-drawing-advertisers-fire.html`, 2012.

[41] Smith, M., Disselkoen, C., Narayan, S., Brown, F., and Stefan, D. Browser history re: visited. In *12th USENIX Workshop on Offensive Technologies (WOOT 18)* (2018), USENIX Association.

[42] Sutton, M. A wolf in sheepâĂŹs clothing, the dangers of persistent web browser storage. *Black Hat DC 2009 Briefings Speakers* (2009).

[43] Swartz, A. Web.py web framework. `http://webpy.org/`, 2018.

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

[45] Van Goethem, T., Joosen, W., and Nikiforakis, N. The clock is still ticking: Timing attacks in the modern web. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security* (2015), ACM, pp. 1382–1393.

[46] Weinber, Z., Chen, E., Jayaraman, P., and Jackson, C. I still know what you visited last summer. In *IEEE Symposium on Security and Privacy. Oakland, CA* (2011), pp. 20–25.

[47] Wondracek, G., Holz, T., Kirda, E., and Kruegel, C. A practical attack to de-anonymize social network users. In *Security and Privacy (SP), 2010 IEEE Symposium on* (2010), IEEE, pp. 223–238.