

When Push Comes to Ads: Measuring the Rise of (Malicious) Push Advertising

Karthika Subramani
University of Georgia
karthika.subramani25@uga.edu

Phani Vadrevu
University of New Orleans
phani@cs.uno.edu

Xingzi Yuan
University of Georgia
x.yuan@uga.edu

Kyu Hyung Lee
University of Georgia
kyuhlee@uga.edu

Omid Setayeshfar
University of Georgia
omid.s@uga.edu

Roberto Perdisci
University of Georgia
Georgia Institute of Technology
perdisci@uga.edu

ABSTRACT

The rapid growth of online advertising has fueled the growth of ad-blocking software, such as new ad-blocking and privacy-oriented browsers or browser extensions. In response, both ad publishers and ad networks are constantly trying to pursue new strategies to keep up their revenues. To this end, ad networks have started to leverage the *Web Push* technology enabled by modern web browsers. As web push notifications (WPNs) are relatively new, their role in ad delivery has not yet been studied in depth. Furthermore, it is unclear to what extent WPN ads are being abused for *malvertising* (i.e., to deliver malicious ads). In this paper, we aim to fill this gap. Specifically, we propose a system called PushAdMiner that is dedicated to (1) *automatically* registering for and collecting a large number of web-based push notifications from publisher websites, (2) finding WPN-based ads among these notifications, and (3) discovering malicious WPN-based ad campaigns.

Using PushAdMiner, we collected and analyzed 21,541 WPN messages by visiting thousands of different websites. Among these, our system identified 572 WPN ad campaigns, for a total of 5,143 WPN-based ads that were pushed by a variety of ad networks. Furthermore, we found that 51% of all WPN ads we collected are malicious, and that traditional ad-blockers and URL filters were mostly unable to block them, thus leaving a significant abuse vector unchecked.

CCS CONCEPTS

• **Information systems** → **Online advertising**; *Web mining*; *Spam detection*; • **Security and privacy** → **Social engineering attacks**.

ACM Reference Format:

Karthika Subramani, Xingzi Yuan, Omid Setayeshfar, Phani Vadrevu, Kyu Hyung Lee, and Roberto Perdisci. 2020. When Push Comes to Ads: Measuring the Rise of (Malicious) Push Advertising. In *ACM Internet Measurement Conference (IMC '20)*, October 27–29, 2020, Virtual Event, USA. ACM, Pittsburgh, PA, USA, 14 pages. <https://doi.org/10.1145/3419394.3423631>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

IMC '20, October 27–29, 2020, Virtual Event, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8138-3/20/10...\$15.00

<https://doi.org/10.1145/3419394.3423631>

1 INTRODUCTION

In the past few years, the rapid growth of online advertising has fueled the growth of ad-blocking software, such as new ad-blocking and privacy-oriented browsers (e.g., Brave [4]) or browser extensions (e.g., AdBlockPlus [1]). In response, both ad publishers and ad networks are constantly trying to pursue new strategies to keep up their revenues. To this end, ad networks have started to leverage the *Web Push* technology enabled by modern web browsers [45]. Until relatively recently, push notifications were mostly limited to native apps on mobile platforms, and web-based applications were unable to connect to their users out of active browsing sessions. However, now Web Push allows for web applications to send out Web Push Notifications (WPN) at any time to re-engage their users, even when the browser tab in which the web application was running is closed (the browser itself needs to be running, but does not need to be in the foreground for a WPN to be delivered to the user). Furthermore, unlike push notifications from native mobile apps, WPNs allow for notifications to be displayed on both desktop and mobile devices. Thus, they serve as a single tool with support to reach users on multiple platforms.

Although WPNs were initially designed for websites to deliver simple messages (e.g., news, weather alerts, etc.), they have become an effective way to also serve online ads, and can therefore be abused to also deliver malicious ads. In particular, the use of WPNs for ad delivery has some unique advantages. First, unlike traditional online ads (banner ads, pop-up ads or pop-under ads), advertisers do not have to wait for users to reach the web page that publishes the ad. Instead, advertisers can send out notifications that can allure users to their targeted content. Secondly, thanks to years of experience with native mobile app notifications, users have been trained to compulsively interact with push notification messages (at least on mobile devices). WPN-based ads may also be less prone to *ad blindness* [60], compared to traditional web ad delivery mechanisms such as page banners. Furthermore, ad-blocking software are not currently effective at blocking WPN-based ads (see Section 6.4), in part because browser extensions are not allowed to interfere with the *Service Workers* code through which WPNs are delivered [48]. For these reasons, some ad networks are focusing their business specifically around WPN ads (e.g., RichPush [17]).

As WPNs are relatively new, their role in ad delivery has not yet been studied in depth. Furthermore, it is unclear to what extent WPN ads are being abused for *malvertising* (i.e., to deliver

malicious ads). In this paper, we aim to fill this gap. Specifically, we propose a system called PushAdMiner that is dedicated to (1) *automatically* registering for and collecting a large number of web-based push notifications from publisher websites, (2) finding WPN-based ads among these notifications, and (3) discovering malicious WPN-based ad campaigns. To build PushAdMiner, we significantly extend the Chromium browser instrumentations developed by [39] and [58], which have been open-sourced by the respective authors. Specifically, neither [39] nor [58] are able to track the activities of Service Workers in detail. Therefore, we implement our own set of browser instrumentations that allows us to track WPNs in all their aspects, from registration to notification delivery, on both desktop and mobile devices. We then build a custom WPN crawler around our instrumented browser to automatically receive, track, and interact with generic WPNs, including collecting malicious WPN ads and their respective malicious landing pages. Finally, we develop a data mining pipeline to analyze the collected WPNs and discover malicious WPN-based campaigns.

To the best of our knowledge, ours is the first systematic study that focuses on automatically collecting and analyzing WPN-based ads and on discovering malicious ad campaigns delivered via WPNs. In contrast, previous work focused on other security-related aspects of Service Workers and Push Notifications, such as building stealthy botnets [48], or social engineering attacks that attempt to force users into subscribing to push notifications [58], but without studying the resulting push messages. Lee et al. [38] study Progressive Web Apps. They collect Service Worker scripts from top-ranked website homepages and analyze their push notifications. Their work studies potential security vulnerabilities related to Service Workers, App Cache, and discusses how push notifications may be abused to launch phishing attacks, without measuring how prevalent these attacks are in the wild. Our work is different, in that we aim to *automatically* collect and analyze WPN-based ads, to discover WPN ad campaigns, and to measure the prevalence of malicious WPN-based ad campaigns in the wild.

In summary, we make the following contributions:

- We present PushAdMiner, a system that enables the automated collection and analysis of online ads delivered via web push notifications (WPNs) on both desktop and mobile devices.
- To track WPNs, we extend a Chromium-based instrumented browser developed in [39, 58] to allow for a detailed analysis of Service Workers, which are at the basis of WPN deliveries. Furthermore, we build a custom WPN crawler around our instrumented browser to collect and automatically interact with WPNs.
- Using PushAdMiner, we collected and analyzed 21,541 WPN messages by visiting thousands of different websites. Among these, our system identified 572 WPN ad campaigns, for a total of 5,143 WPN-based ads that were pushed by a variety of ad networks. Furthermore, we found that 51% of all WPN ads we collected are malicious, and that traditional ad-blockers and URL filters were mostly unable to block them, thus leaving a significant abuse vector unchecked.

2 MOTIVATING EXAMPLE AND BACKGROUND

In this section, we provide an example of WPN-based malicious ad, and then briefly explain the concepts and technologies behind web notification services.

2.1 Motivating Example

Figure 1 provides an example of malicious WPN-based ad. During the preliminary stages of our research, we stumbled upon a website on aurolog[.]ru. When visiting the main page, the site requested permission to send us notifications. We granted permission by pressing the *Allow* button on the browser dialog box, and subsequently received a WPN ad with the following alert message: “Your payment info has been leaked” (see Figure 1). After clicking on the notification, we were redirected to a *tech support scam* [46]. To our surprise, the landing URL was neither blocklisted by Google Safe Browsing[11] nor detected as malicious by any of the web page scanners on Virus Total[21]. This example confirmed our suspicion that WPNs may be abused for malvertising, and sparked our investigation to determine whether such cases of malicious WPN-based ads could be automatically collected and analyzed.

2.2 Technical Background

Recent changes in HTML5 have introduced new web features, such as *Service Workers*[20], *Push Notifications*[14] and *AppCache*[19]. Websites that adopt these technologies are called Progressive Web Apps (PWAs). Throughout this paper, we refer to push notifications sent by PWAs using a browser as Web Push Notifications (WPN), to distinguish them from push notifications sent by native apps on mobile devices, and refer to Service Workers as SWs for brevity.

Service Workers and Push Notifications: A *Service Worker* (SW) is an event-driven script executed by the browser in the background, separately from the main browser thread and independently of the web application from which it was initially registered and that it controls. In practice, a SW comes in the form of a JavaScript file that is registered against the origin and path of the web page to which it is associated (only HTTPS origins are allowed to register a SW). In effect, SW can be viewed as “a programmable network proxy that lets you control how network requests from your page are handled”[13].

Service Worker can use the *Push API*[16] to receive messages from a server, even while the associated web application is not running. It is worth noting that a single web app is allowed to register multiple SWs. Service Workers can also use the *Notifications API*[14] to display system notifications to the user. A prerequisite is that the web application must first request permission to display notifications to the user (only allowed for HTTPS origins). If the user accepts (i.e., clicks on “Allow” instead of “Block” on the notification request popup) to receive notifications from the web application’s origin, this permission persists across browser restarts, and until the user explicitly revokes the permission via browser settings/preferences (notice that non-expert users may find it difficult to understand, find, and disable notifications in the browser’s settings).

Web notification messages have a number of customizable parameters, such as title, body, target URL, icon image, display image and action buttons. The user can interact with a notification by either

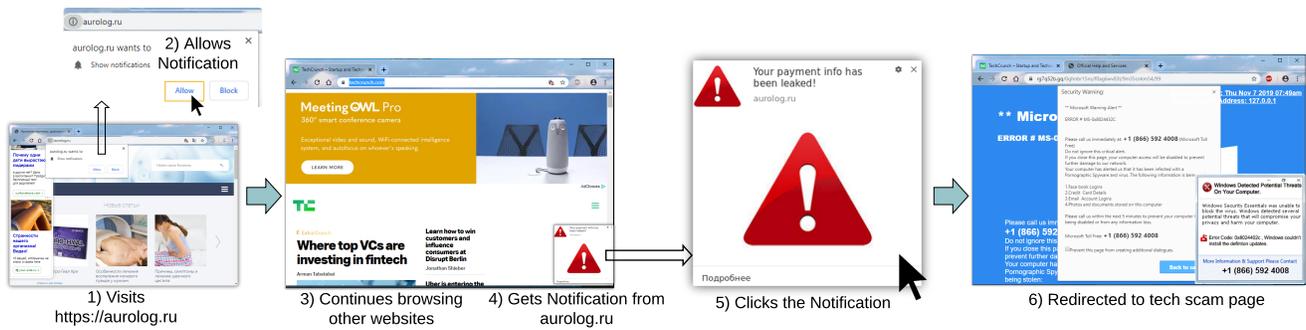


Figure 1: Example of malicious advertisement served through web push notifications

clicking on it, closing it or performing any custom actions displayed in the notification message. SW can listen to such user events and take action according to the input. This includes loading target URL on a separate tab, following a user's click on the notification box.

Firestore Cloud Messaging (FCM): FCM is a cross-platform messaging solution for Push Notifications. It can serve as a central authority that mediates the communication between the ad server and the Service Worker. Upon initial registration, FCM creates a unique registration ID per user and per Service Worker, which is sent along with an endpoint URL [12] to the ad server. For further details, refer to FCM's online documentation [33].

3 SYSTEM OVERVIEW

In this section, we provide an overview of how PushAdMiner works, leaving a detailed description of the main system's components to Sections 4 and 5. A high-level representation of the system is provided in Figure 2.

PushAdMiner consists of three main components: (i) an instrumented browser to collect fine-grained information about SWs and WPNs; (ii) a custom crawler that automatically visits sites and interacts with the browser, including granting notification permissions and interacting with WPNs (Section 4); and (iii) a data analysis component aimed at identifying WPN-based ad campaigns and labeling likely malicious ones (Section 5).

While a number of browser automation and crawling systems have been proposed, including Selenium [53], Puppeteer [32], and others [36, 39, 58], currently they do not fully support the automatic user interactions with WPNs and collection of all details about SWs needed for our study. We therefore built an instrumented browser based on Google Chromium, by significantly extending existing open-source Chromium instrumentations [39, 58]. In addition, we leveraged Puppeteer [32] for browser automation and event logging, and wrote custom scripts to record SW registrations and network requests. Figure 2 presents an overview of how PushAdMiner collects information about WPNs, and how the browser logs are analyzed to identify ad campaigns in general and discover malicious ones among them. First, we discover a set of URLs that may send push notifications with the help of an ad network and filter those that actually request for notification permission (see Section 6.1 for details). For the web pages that ask for notification permissions, we log details about the responsible SW code, automatically grant

permission (via browser code instrumentation), and then collect the notifications that are later pushed to our instrumented browser. When a notification is displayed by the browser, we record fine-grained details about the notification message itself (including message text and icons), automatically simulate a user click on the notification box (via browser code instrumentation), and track all events resulting from the click. If the click results in a new page being open, we record detailed information about the related network requests, including all browser redirections, as well as detailed logs and a screenshot of each new page the browser visits, including the landing page (i.e., the final web page reached due to the click).

Finally, we extract relevant information from the detailed logs of our instrumented browser, and apply a clustering strategy to find notifications that are similar to each other, which allows us to identify WPN-based ad campaigns. We then leverage URL blocklists to find WPN ad campaigns that are likely malicious (e.g., because one or more landing pages are known to be malicious).

Note that in this paper we do not focus on building a malicious WPN ad campaign detector, such as using statistical features or machine learning classifiers. Rather, our focus is on discovering, collecting, and analyzing WPN ad campaigns in general, and on measuring the prevalence of both benign and malicious campaigns. As we will show in Section 6, URL blocklists tend to miss a significant number of malicious URLs that we determine to be related to malicious ad campaigns. The analysis we present in this paper could therefore be used as a starting point for developing an automated malicious WPN ad campaign detector. We leave this latter task to future work. Our code to collect and analyze WPNs is publicly available in a Github repo¹ and a Docker container with the instrumented Chromium is in Docker Hub².

Ethical Considerations: To track WPN-based ads and label malicious ones, it is necessary to collect information about the landing page that an ad eventually redirects to. For instance, for most malicious ads the attack is effectively realized only once the user reaches the landing page, especially in case of social engineering and phishing attacks. As we do not know in advance what landing page will be reached by clicking on a WPN message, and whether a WPN ad is malicious or not, our system will likely click on both

¹<https://github.com/karthikaS03/PushAdMiner>

²https://hub.docker.com/repository/docker/dockerammu/docker_puppeteer_chromium_xvfb

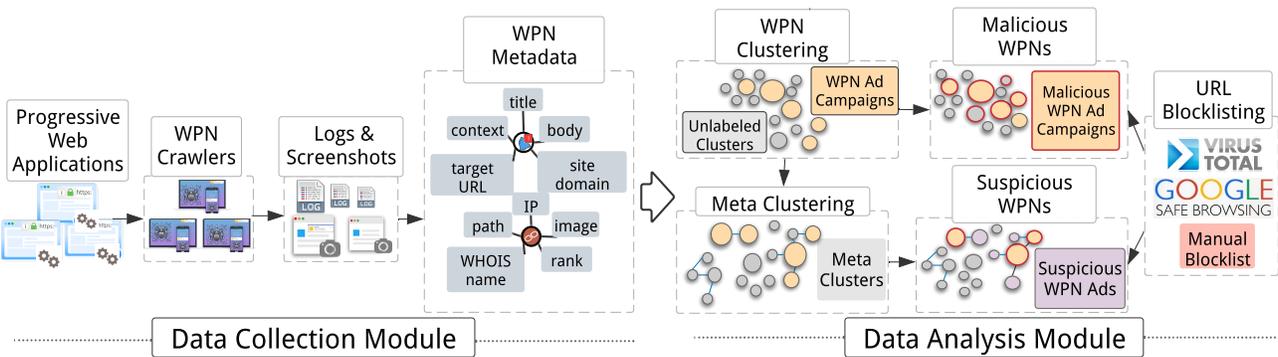


Figure 2: PushAdMiner System Overview

legitimate and malicious ads. In turn, this may cause legitimate advertisers to incur a small cost for our clicks, as they will likely have to pay a third-party publishing web page and ad network for their services (notice that we obviously receive no monetary gain whatsoever during this process). This is common to other similar studies, such as [49, 58], and we therefore address the ethical considerations for our study following previous work.

To make sure we do not have a significant negative impact on legitimate third-parties, we estimated the cost incurred by these advertisers due to ad clicks performed by our system, and found that our system has negligible impact on advertisers. Specifically, among the WPN ads we identified, we consider legitimate ads to be those whose landing pages are not labeled as malicious by Virus Total’s URL classification services. Then, we estimate the *cost per landing domain* based on the number of ads we clicked on that lead to a specific domain, using the *Cost Per Mille (CPM)* [5] for push notification ads according to iZooto [18]. The maximum cost per landing domain throughout our entire study was USD 1.12 (due to landing on the same domain 444 times), which we calculated using the standard CPM of USD 2.54. On average, we visited each landing domain 18 times, which corresponds to an average cost of USD 0.04 per landing domain (i.e., per advertiser). Considering these low values, we believe the impact of our system on advertisers is not significant, and is on par with previous work [49, 58].

4 DATA COLLECTION MODULE

In this section, we describe in detail how PushAdMiner’s data collection module is implemented. The steps referred to in the following sections follow the numbering given in Figure 3.

4.1 Desktop Environment

To discover WPN ad campaigns, we first need to collect WPN messages. To this end, we build a crawler consisting of an instrumented browser and browser automation scripts. As our crawler encounters a website that asks for permission to send push notifications, our goal is to automatically allow the permission request, so that we can collect notifications from that origin. To this end, we instrumented our browser as follows: we introduce changes to the RequestPermission and PermissionDecided methods under the PermissionContextBase class in Chromium’s C++ code base, to

log all details about the permission request (e.g., the origin requesting it) and to automatically grant permissions.

Service Worker Registration (step 2): To record when a SW is registered by a web page, we use a custom script for Puppeteer [32]. Specifically, we listen to serviceworkercreated events and log information such as details about the page that registered the SW, and the URL from which the SW code was retrieved.

Network Requests (step 3): When the browser receives a push message from the FCM service, it will alert the SW to which the message is destined. The correct SW is identified via a unique ID included in the push message. As part of handling the push message and related notification, the SW may issue one or more network requests directly to an ad network server or other third-party servers. For instance, the SW may contact an ad network server to determine the landing URL associated to a given push message. Also, after a pushed notification receives a user click, the SW may send a network request to an ad server to notify it of the user click, or to facilitate other tracking related activities. We capture and log such network requests issued by SWs by leveraging a custom Puppeteer [32] script that listens to SW request and response events. Specifically, for every SW’s network request we record information such as the requested URL, the data sent/received, possible network redirections, and the related response content.

Notifications (steps 5-6): Notifications are displayed by invoking showNotification under ServiceWorkerRegistrationNotifications (in Chromium’s C++ code). We therefore add an instrumentation hook to record calls to that method. Specifically, we log the URL of the SW that called for showing the notification, the notification title, body, icon URL, and target URL (if present).

Notification Clicks and Navigations (steps 7-8): Depending on the type of notification, clicking on it can navigate the browser to a new web page (e.g., on a separate tab). For WPN-based ads, clicking on the notification box typically takes the user to the page advertised in the WPN ad (i.e., the ad’s landing page). Collecting such pages is especially important for studying malicious WPN ads, because the landing page often implements a critical component of the attack.

For instance, Figure 1 shows that the user reaches the actual *tech support scam* page after clicking on the notification. The landing

page advertises the scam phone number, without which the attack could not be monetized.

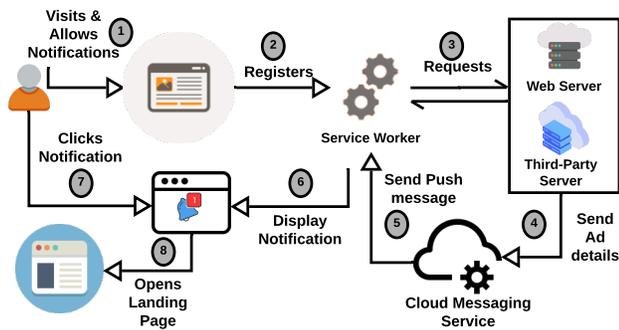


Figure 3: Steps involved in Serving Ads via WPNs

To automate the process of collecting the landing page associated to push notification messages, we need to simulate a user click. Unfortunately, browser automation frameworks such as Puppeteer [32] and Selenium [53] do not allow for interacting with WPNs. Therefore, we again had to build our own custom browser instrumentation. Specifically, we found that Chromium uses the `Add` method in `MessageCenterNotificationManager` (in the C++ code base) to trigger the display of a notification. Also, we found that `WebNotificationDelegate` has a method called `Click` that is responsible for propagating user clicks to the notification. Therefore, to simulate a user click we instrument the `Add` method mentioned above so that, after the notification is displayed, it waits for a short delay (e.g., a few seconds) and then calls `WebNotificationDelegate::Click`. If the `Click` call results in a web page navigation, our browser records all network requests involved in the process (including all redirections) and fine-grained details about the rendering of the landing page, which includes detailed information about any JavaScript code executed in the context of the landing page.

4.2 Mobile Environment

We also developed a version of `PushAdMiner` for Android. Due to some technical differences between how WPNs are displayed on a mobile OS, compared to desktop environments, we had to adapt some of the system components to run specifically on Android. First, at the time when we started building our system, `Puppeteer` [32] did not appear to support Android Chromium automation. Only recently there have been online posts in which `Puppeteer` users describe how they have been able to “hack” their configurations to remotely control an Android browser. We therefore built our own browser automation framework that works via the `Android Debug Bridge (ADB)`. The capabilities of our `ADB`-based automation framework are limited, but sufficient for enabling data collection for `PushAdMiner`. We plan to explore the use of `Puppeteer` for `PushAdMiner` on Android as future work.

Logging Internal Browser Events: We compile our instrumented Chromium browser for Android, so that we can collect intimate details about internal browser events related to WPNs, including recording information about the related `Service Workers` and the

rendering of the landing page resulting from clicking on a WPN. Browser logs are sent via the `logcat ADB` command to a remote logging machine.

Interacting with Notifications: Unlike on desktop devices, in which WPN messages are displayed by the browser, on Android device it is the Android OS that displays a WPN as a system notification. Also, unlike on desktop environments, the browser does not need to be activated for a WPN message to be received, though the browser may be activated after tapping on a notification (e.g., to navigate to the URL pointed to by the notification). We therefore had to implement a different mechanism to simulate user interactions with WPNs on Android. Specifically, we developed an Android application that leverages Android’s `Accessibility Service`. The `Accessibility Service` is aimed to help people with disabilities in using the device and apps. It is a long-running privileged system service that helps users process information from the screen and lets them interact with the content meaningfully in an easy way. Android developers can leverage the `Accessibility Service API` and develop apps that are made aware of certain events, such as `TYPE_VIEW_FOCUSED` and `TYPE_NOTIFICATION_STATE_CHANGED`. Furthermore, the `accessibility service API` can also be used to initiate user actions such as click, touch and swipe.

We install our app with `Accessibility Service` permission on an Android physical device, and use it to interact with every notification event fired. Whenever a new notification pops up, our application will automatically swipe down the notification bar and click on the notification to complete the action, while our instrumented Android browser produces detailed logs about the consequences of such interactions (e.g., loading a new web page).

As an alternative, `AndroidDebugBridge (ADB)` could be leveraged to implement the same browser automation that we implemented using `Accessibility Service`. However, in practice, we found that it could create large traffic overhead through the USB cable connected with the device. As we already use `ADB` to retrieve fine-grained browsing logs, we decided to avoid further USB overhead. Therefore, we found the use of the `accessibility` features to be better in practice for our specific application.

5 DATA ANALYSIS MODULE

In this section, we first describe in detail how we mine the collected WPN messages to identify WPN-based ad campaigns, and later explain how we label malicious campaigns.

5.1 WPN-based Ad Campaigns

To identify WPN-based ad campaigns, we mine the dataset of WPN messages collected by `PushAdMiner` from a large and diverse set of websites. To distinguish between generic WPN messages and WPN-based ads, we start by considering the following intuitions. Roughly speaking, advertisers tend to promote their products or services on multiple websites, and clicking on a WPN notification typically leads to a third-party landing page, on a different origin than the website from which the WPN message was received. On the other hand, non-ad WPN messages are typically related to alerts (e.g., breaking news, weather alerts, etc.) that are specific to the notifying website itself, and clicking on the notification often leads to the same origin to which the notification was related.

Following the above intuitions, we broadly define a *WPN ad campaign* as a set of WPN messages from multiple sources that deliver similar content promoting the same (or similar) products or services. In practice, this translates into a (potentially large) group of similar WPN messages pushed by multiple different websites that lead to the same landing page, or different landing pages that show similar content.

To find such WPN ad campaigns among a large collection of generic WPN messages, we leverage a document clustering approach. This clustering process aims to group together WPN messages that are similar to each other in terms of message content and landing page. As a result, WPN ads that belong to a large WPN ad campaign will tend to form larger clusters. On the other hand, non-ad WPNs would tend to be isolated into singleton clusters or clusters that contain messages related to only one source website and a landing domain that points back to the source website itself.

Because clustering is an unsupervised learning process, it is usually challenging to tune the hyperparameters to obtain perfect clustering results. Therefore, to minimize the chances of grouping together ads and non-ads, we tune our clustering system to be conservative and yield “tight” clusters. Namely, WPN messages assigned to the same cluster will exhibit high similarity. Although this may result in some WPN ads being left out of the campaign (i.e., the cluster) they may also belong to, we will see later that we can reconnect them to their respective campaigns using a meta-clustering step (Section 5.3).

5.1.1 WPN Clustering Features and Approach. To cluster similar WPNs, we first need to define the features to be compared and a similarity function that calculates the closeness of two WPN messages. Thanks to our instrumented browser, each WPN we collect is accompanied by metadata such as the notification title, message body, the images and icons found in the WPN, the URL of the landing page, a screenshot of the landing page reached after clicking on the WPN (in the case of our desktop browser), the events that occur immediately after a notification click, etc.

Among the above features, to discover WPN ad campaigns we use the following ones for clustering: message title, message body, and the URL path (i.e., we exclude the domain name) of the message’s landing page URL. We then use these features to calculate the pairwise distance between WPN messages. Specifically, given two messages, we compute the distance between their title and body text, and separately the distance between their landing URL paths. Then, we compute the total distance between two WPNs as the average of these two distances. We explain the distance measures in more detail below. The remaining information contained in the WPN metadata collected by our browser (e.g., domain names, screenshots, etc.) that are not used as features will instead be used later to help us validate the clustering results.

WPN Messages Distance: The text contained in the title and body of a WPN message is typically short and includes specific keywords that reflect the *theme* of the message. To measure the similarity between these short pieces of text (the concatenation of title and body), we require a measure that gives importance to significant keywords found in the content of WPNs. To this end, we use the *soft cosine similarity* [54] measure, which considers the semantic relationship between words. To compute the similarity between words, we first use

Word2Vec [8] on all WPN messages to obtain a term-similarity matrix. We then convert each WPN message into a bag-of-words representation, and input both the term-similarity matrix and the bag-of-words vectors for each pair of WPN messages into the cosine similarity function (we use `softcosim()` implemented in `gensim`[31]) to obtain a pairwise similarity matrix for WPN messages. Since the similarity value $s \in [0,1]$, we calculate the distance as $d = (1-s)$.

URL Paths Distance: Given the landing page URL of a WPN (which we collect along with other metadata, as explained earlier), we extract the URL path by excluding the domain name and the query string values, while retaining the relative path to the page and two query parameter names. To calculate the distance between two URL paths, we use the Jaccard distance between tokens extracted from the URL path, such as the components of the directory path, the page name, and the name of the query string parameters.

WPN Clustering: Once the pairwise distances are computed, we use agglomerative hierarchical clustering over the distance matrix to form *WPN clusters*. To determine where to cut the resulting dendrogram, we compute the average silhouette score [52] for clustering results obtained at different cuts, and choose the cut with the highest score.

WPN Ad Campaigns: As per our definition of WPN ad campaigns, to determine if a WPN cluster is formed by ads, we take into consideration the source of the WPNs. Namely, we count the number of effective second-level domain names associated with the websites that sent the push notifications per each cluster. This number tells us if the WPNs have been published on multiple sources. If a cluster contains more than one distinct second-level source website domain, we label the cluster as a *WPN ad campaign*.

5.2 Identifying Malicious WPN Clusters

To determine the maliciousness of a WPN cluster, we leverage two well-known URL blocklisting services: Google Safe Browsing (GSB)[11] and Virus Total[21] (VT). We submit the full URLs of all the landing pages reached from all WPN messages in the cluster to these services. Then, we label a particular WPN message as *known malicious* if the landing page URL is blocklisted as malicious by any of the blocklisting services. Next, we use a simple label propagation policy to flag as malicious the WPN clusters that contain at least one known malicious WPN.

It is worth noting that we submit full URLs to GSB and VT. If a full URL, u , is blocklisted, it does not necessarily mean that all URLs under the same domain name as u will also be blocklisted (in fact, we found some cases that confirm this observation for VT). In addition, similar malicious WPN messages often lead to different domain names, mainly as an attempt to evade blocking by URL blocklists. At the same time, because WPN messages in the same cluster are very similar, thanks to our conservative clustering approach, they share very similar title, body, and structure of the landing page URL path. Intuitively, if one WPN message is known to lead to a malicious landing page, it is highly likely that all other WPN messages in the same cluster also lead to similar malicious pages, as we will also discuss in our measurement results (see Section 6). That is why we rely on the simple “guilty by association” policy mentioned above for label propagation.

5.3 WPN Meta Clustering

Because URL blocklists have limited coverage, it is possible that some malicious WPN clusters will not be immediately identified using the labeling approach discussed above. Furthermore, because our clustering approach is conservative, it is possible that separate clusters of WPN messages may in fact be related to each other. To compensate for this, we perform a meta-clustering step that aims to group together WPN clusters that may belong to the same WPN “operation” (e.g., the same advertiser), as explained below.

To this end, we generate a bipartite graph $\mathcal{G} = (W, D, E)$. W is a set of nodes in which each node represents a WPN cluster obtained as described in Section 5.1.1. D is the set of all domains pointed to by the WPN messages we collected (i.e., all domains related to any of the landing page URLs found in the WPN messages), and E is a set of edges in which each edge connects a node $w \in W$ to a node $d \in D$. Specifically, we connect each WPN cluster to the domains related to the landing page URLs linked by the WPN messages in the cluster. Then, we find all isolated components, $G_1, G_2, \dots, G_m \in \mathcal{G}$, and consider each isolated component as a *meta cluster* of WPN messages. Notice that this leads us to groups of WPN clusters that are related to each other because they collectively share common landing page domains. Figure 5 visually shows two examples of such meta clusters.

5.4 Identifying Suspicious WPN Clusters

In this section, we provide a systematic way to identify “potentially malicious” (also referred to as suspicious) WPN ads in addition to the malicious WPN ads that were found in the previous section. Further, this helps to drastically reduce the manual effort required to analyze all WPNs.

First, we consider those WPNs as *suspicious* whose landing domain names are associated with full URLs that have been previously labeled as malicious by GSB or VT. Let $G_i \in \mathcal{G}$ be a meta cluster that includes one or more of such known malicious URLs. Then, we label WPN clusters in G_i (and thus all WPN messages in those clusters) as suspicious, unless they were previously labeled as malicious according to the process described in Section 5.2.

In addition, given a meta cluster G_i , if it contains at least one WPN cluster $w_j \in W$ that we previously identified as a *WPN ad campaign* (see Section 5.1), we consider all WPN messages contained in the WPN clusters within G_i as WPN ads. This is because those WPN messages point to domain names related to WPN-based advertising, since they are linked to one or more WPN ad clusters, and thus are highly likely WPN-based ad messages themselves.

Apart from using their association with malicious URLs, to identify additional suspicious WPN ad campaigns that were not previously labeled based on the process described above and in Section 5.2, we proceed as follows. We notice that ad networks such as Google Ads and Bing Ads recommend advertisers not to promote the same product, brand, or similar customer experience by redirecting users to multiple destination websites. Violating this policy is referred to as *Abuse of Ad Network* by Google Ads[9] and *Duplicate Ads* by Bing Ads[3]. Besides the fact that these practices do not conform to advertising policies, malicious advertisers often prefer hosting similar malicious content on multiple domains to evade detection and to continue the attack even if one of their domains is blocklisted. We then identify such practices and label the related

meta-clusters that lead to multiple different landing domains as *suspicious*, to trigger further (manual) analysis to confirm if they are malicious. We provide detailed measurements on the identified suspicious WPN ads in Section 6.3.3.

Manual Verification: To validate the malicious URLs detected by URL blocklists, and to measure the number of suspicious WPN clusters that are in fact malicious, we manually analyze all malicious and suspicious WPN clusters discovered by our data analysis module. During manual analysis, we consider multiple factors to determine if the WPNs are indeed malicious. Once we manually confirm that a WPN cluster is malicious, we add it to a *manual blocklist*, which we then use to inform the measurement results presented in Section 6.

Some of the factors considered during manual blocklisting are as follows. We recognize a WPN message as malicious if it meets a combination of the following conditions: (1) leads to a landing page visually similar to a known malicious page (as determined by GSB and VT); (2) contains the same WPN message as a known malicious WPN message, but leads to a different product/site (i.e., a different landing page); (3) includes message content that is highly likely malicious or leads to a page with likely malicious content such as rewards that clearly seem too good to be true, as is typical of survey scams [37], false financial alerts, etc.; or (4) leads to a landing page that shares several domain-related properties with known malicious sites, such as IP address, registrant, similar domain name, etc. We further discuss our manual analysis process by presenting some examples in Section 6.3.2.

6 MEASURING WPN ADS IN THE WILD

In this section, we report measurements on the usage of WPNs as an ad delivery platform, and provide insights into the malicious use of WPN ads.

6.1 Data Collection Setup

We first describe PushAdMiner’s setup for harvesting in-the-wild WPN messages for both desktop and mobile environments. Because our internal browser instrumentations are implemented by extending the browser code provided by [58], our data collection process leverages Chromium’s code base version 64.0.3282.204, which we built for both Linux and Android environments.

6.1.1 Seeding the WPN Crawler. Our main goal is to setup the WPN data collection system to maximize our chances of collecting WPN-based ads, so that we can measure their properties and discover abuse. To this end, we rely on the publicwww.com code search engine. Compared to simply crawling popular websites, this allowed us to more efficiently discover a diverse set of URLs (both popular and unpopular) that use push notifications. We first identify a *seed* set of popular advertisement networks that support push notification advertisements. Specifically, by manually performing online searches and reading articles and forums dedicated to advertising online, we manually discovered 15 popular ad networks that provide push advertisement services. We registered an account with 9 of these ad networks to obtain the JavaScript code that needs to be embedded in an ad-publishing websites to include the ad networks’ SWs that will control their push notifications requests. We then used the publicwww.com code search engine to find websites (i.e., URLs) that likely embed the identified ad network’s code. The code

search engine (publicwww.com) itself provides a list of *ad networks to search keywords* mappings [15], which we leveraged to obtain URLs of websites that likely embedded ads from the remaining 6 seed ad networks of interest, without the need to register an account with those networks as well. The list of 15 seed ad networks we discovered this way is shown in Table 1.

In addition to the search keywords related to the 15 ad networks mentioned above, we further used *generic keywords* that potentially indicate that a web page employs push notifications, such as `NotificationrequestPermission`, `pushmanagersubscribe`, a 'push' parameter value for `addEventListener`, etc. Overall we obtained a list of 19 code search keywords to be used for searches on publicwww.com that are likely to lead to URLs that make use of push notifications. These keywords allow us to explore a large set of web pages that goes beyond websites where ads from the seed ad networks are most likely hosted. Table 1 provides a count of the HTTPS URLs found via publicwww.com for each ad network and generic WPN-related keyword.

As a result of the above search, we were able to gather a total of 87,622 HTTPS URLs that could potentially register service workers and publish WPN ads (notice that SWs can only be registered by HTTPS origins). These URLs were hosted on 82,566 distinct second-level domain names. We use this set of URLs as seed for PushAdMiner’s WPN message crawlers. Notice that there is no guarantee that the sites we crawl will in fact lead us to registering for and receiving WPN-based ads. Rather, only a subset of these web pages will actually request notification permissions. Therefore, we visited each URL and retained only those that actually make a request for a notification permission. Then, we used PushAdMiner to automatically grant notification permission requests on those URLs.

As shown in the last column of Table 1, overall we identified 5,849 URLs hosted on 5,697 distinct second-level domains that issued a notification permission request. Of these 5,697 domains, 2,040 (36%) ranked within the Alexa.com top one million most popular websites, and 1,383 of them actually sent us notifications during our experiment period. Table 2 shows a detailed breakdown of how the 2,040 domains ranked among the Alexa.com most popular websites. These results show that our process for collecting WPNs is not limited to unpopular, “low-tier” websites. In addition, they show that even highly ranked websites can push WPN ads, some of which resulted to be malicious.

6.1.2 Collecting WPNs in a Desktop Environment. To automatically harvest WPN advertisements at a large scale, we leverage Docker containers [29] to launch several parallel instances of our instrumented browser-based crawlers. During our experiments, we used four different Ubuntu 16.04 Linux machines having between 8 to 32 CPU cores and 64 to 128 GB of memory each, running a total of 20 to 50 Docker sessions in parallel at a time. For each seed URL that issues a notification permission request, we start the monitoring phase described below.

Every time we visit a URL, we wait 5 minutes to make sure the website has sufficient time to present our browser with a notification request. Once a permission request is received and automatically granted by a browser instance, a SW is registered. If a SW was registered, we keep the related Docker container alive for an additional 15 minutes, to allow the browser to receive the first (or more than

one) WPN message from the visited URL. To select this 15 minutes threshold, we first performed pilot experiments with much longer waiting times (up to 96 hours) for a large subset of URLs (1,425 URLs, to be precise), and observed that 98% of them sent their first notification within 15 minutes of when the permission was first granted.

Given a container, and therefore a browser instance that granted notification permissions to a specific URL, after the first 15 minutes of its life we suspend the container to free up resources for instantiating a new container that will visit a new URL. However, we periodically resume suspended containers to see if they will receive additional notifications, which are queued in the FCM and sent to the browser as it comes back online.

Table 1: URLs and Notification Permission Request counts

Ad Network	URLs	NPRs
Ad-Maven	49,769	1,168
PushCrew	15,177	427
OneSignal	11,317	2,933
PopAds	1,582	73
PushEngage	796	215
iZooto	676	278
PubMatic	647	7
PropellerAds	335	9
Criteo	154	5
AdsTerra	115	2
AirPush	52	0
HillTopAds	21	3
RichPush	12	0
AdCash	10	0
PushMonetization	9	5
Generic Keywords	URLs	NPRs
NotificationrequestPermission	3,965	538
pushmanagersubscribe	2,667	158
addEventListener('Push')	263	9
adsblockpushcom	55	19
Total	87,622	5,849

Table 2: URL rankings on Alexa’s top 1M sites

Alexa Ranking	Request Notifications	Sent Notifications	WPN Ads (# of Sites)	Malicious WPN Ads (# of Sites)
< 5k	37	21	6	3
< 10k	93	46	8	5
< 50k	317	151	21	11
< 100k	515	230	33	14
< 500k	1494	610	90	33
< 1M	2040	1383	135	60

6.1.3 Collecting WPNs in Mobile Environment. During our study, we found that WPN messages sent to mobile devices tended to be somewhat different than the ones collected by desktop browsers, in that they were more tailored to mobile users. In particular, malicious mobile WPN messages included fake *missed call* notifications, fake *amber alerts*, “spoofed” *Gmail* or *WhatsApp* notifications, fake *FedEx* notifications, etc. In addition, we found that these malicious messages were much more likely to appear on real Android devices, rather than emulated environments (likely due to some form of emulator detection). Therefore, to automatically collect mobile WPN messages we instrumented a real mobile device. Specifically, we used a Google Nexus 5 device with 2 GB of RAM and a 1080×1920 pixels display. The Android version we used was `aosp_shamu-userdebug7.1.1N6F26Y`.

As we attempted to scale our PushAdMiner’s mobile WPN crawlers on a real device, we identified two challenges. First, Docker or other

container techniques do not support Android, and therefore we cannot easily visit multiple URLs in parallel with isolated browsing sessions. Second, we considered to use app cloning techniques [57] to open multiple browser instances separately in isolated execution environments. However, the limited computing power of mobile device restricted us to scale up and visit a large number of URLs simultaneously. Therefore, we decided to open multiple URLs in one chromium app but in separate tabs.

Table 3: Measurement Results of Data analysis Module

	WPNs with Landing Pages	WPN Ad Campaigns	WPN Ads	Malicious WPN Ad Campaigns	Malicious WPN Ads
Desktop	9,570	572	5143	318	2615
Mobile	2,692				
Total	12,262				

6.2 WPN Messages Dataset

We start with the 5,849 initial URLs that we collected as explained in Section 6.1.1, over 5,697 distinct second-level domain names. By clicking on WPN messages issued by these initial URLs, we collect an additional 10,898 URLs across 2,269 distinct second-level domains. When visited, many of these additional URLs presented our browser with a notification request, which our crawler automatically granted. This brought us to a total of 7,951 URLs that registered a SW with Push permission and were therefore able to push notifications to our instrumented browser instances over time.

During the course of about two months (September and October 2019), we were able to collect a total of 21,541 push notification messages, including 12,441 notifications for the desktop environment and 9,100 for the mobile environment. PushAdMiner interacted with each of these WPN messages. However, not all automated clicks on notification boxes led to a separate landing page. In addition, some landing pages appeared to cause a crash in the browser’s tab (but not the browser) in which they rendered, preventing us from collecting detailed information on those pages (this was likely due to the fact that our instrumented Chromium browser is not based on the most recent stable code base). We filtered out these notifications, leaving us with 12,262 WPN messages (9,570 on desktop and 2,692 on mobile) that when clicked on lead to a valid landing page. We then used this final set of WPN messages for clustering process (refer Section 5).

6.3 Data Analysis Results

Summary of findings: Table 3 summarizes the overall results of our analysis process. From the 12,262 WPN messages mentioned above, PushAdMiner identified 572 WPN ad campaigns and a total of 5,143 WPN ads related to these campaigns. Moreover, PushAdMiner identified 51% of all WPN ads as malicious. Specifically, PushAdMiner found 318 (out of 572) campaigns to be malicious; in aggregate, these malicious campaigns included 2,615 WPN ads.

This is quite a staggering result, in that it appears that ad networks that provide WPN ad services are heavily abused to distribute malicious content. Later, in Section 6.4, we also show that ad blockers are ineffective at blocking such ads, which is an additional cause

of concern. In the following sections, we discuss the clustering and labeling results in more detail.

6.3.1 WPN Clusters and Ad Campaigns. As discussed in section 5.1.1, we cluster the collected WPN messages based on their message content and landing page information. After clustering 12,262 WPN messages that led to a valid landing page, we obtained 8,780 WPN clusters, of which 7,731 were singleton clusters containing only one element (i.e., only one WPN message). Of the remaining non-singleton clusters, 572 were labeled as WPN ad campaigns, according to the process described in Section 5. In aggregate, these WPN ad campaigns pushed 3,213 WPN ad messages to our browsers, during a period of about two months.

Figure 4 provides some concrete examples of WPN clusters. In both WPN-C1 and WPN-C2, the respective WPNs were pushed from multiple sources (i.e., multiple second-level domain names), as also shown in Figure 4. WPN-C3 included 4 identical WPN messages pushed by a single source website, a bank, alerting users on their loan offers. These messages appear to be legitimate, and led back to the site that pushed them. WPN-C4 is an example of WPN message isolated into a singleton cluster. According to the definition provided in Section 5, we label WPN-C1 and WPN-C2 as WPN ad campaigns, because the WPNs in each of the clusters deliver very similar (or the same) message promoting very similar products from multiple sources. However, WPN-C3 and WPN-C4 do not meet the definition and are thus not labeled as WPN ad campaign.

6.3.2 Malicious WPN Ad Campaigns. As described in Section 5.2, we submit landing page URLs related to all WPN messages to GSB [11] and VT [21]. On our initial scan, less than 1% of the URLs were detected as malicious by GSB or VT, in aggregate. For instance, initially VT flagged 108 landing page URLs as malicious, of which 88 were related to WPN messages labeled by our system as belonging to ad campaigns. Notice that for VT we consider a URL as malicious if at least one of the URL detection engine reports it as malicious, and later manually review all results to filter out possible false positives. After one month, we submitted the same set of URLs once again, and we found that 1,388 URLs (11.31%) were detected by VT, though GSB still only flagged 1% of them.

Table 4: Measurement Results at Stages of Clustering

	# clusters	# ad-related clusters	# WPN ads	# known malicious ads	# additional malicious ads
After WPN Clustering	8780	572	3213	758	367
After Meta Clustering	2046	224	1930	210	1280
Total:		5143		968	1647

PushAdMiner relies on label propagation to label WPN messages and clusters as malicious, based on results from VT and GSB, as explained in Section 5.2. To limit the chances of amplifying possible false positives from VT and GSB, we manually verified all 1,388 URLs to check whether they actually led to malicious content. We were able to confirm that 96.8% of them indeed appeared to be malicious. Of the remaining 44 URLs that we could not confirm as malicious, 13 were found to belong to popular benign domains such as bing.com, kbb.com, tophatter.com, etc.; 24 URLs were related

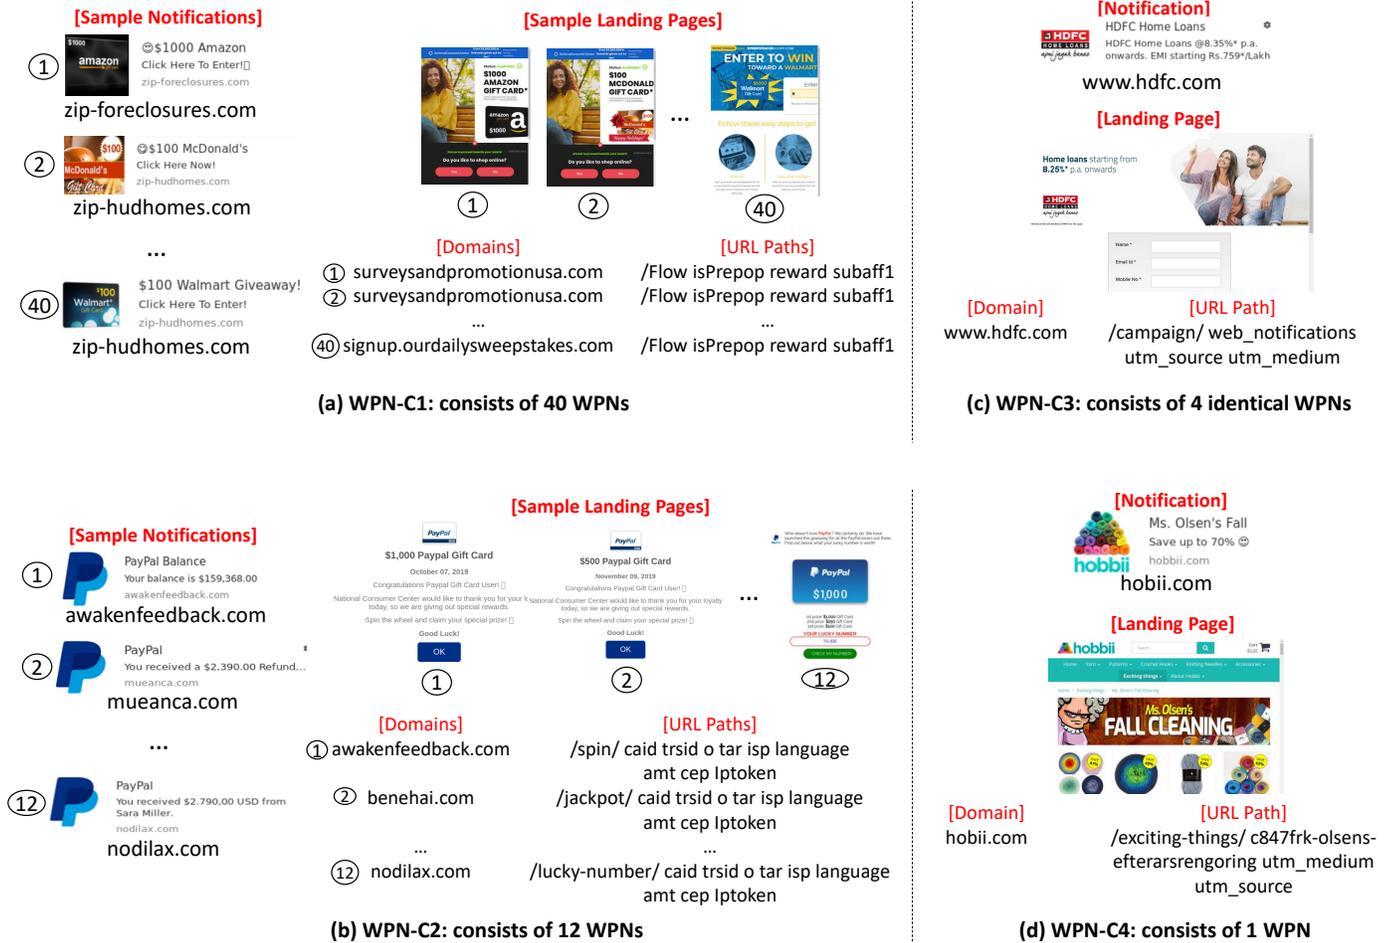


Figure 4: Examples of WPN clusters

to unpopular blog/news sites; 3 led to adult websites; and 4 led to websites hosting non-English content that we could not verify. Given that these sites may be benign, since we do not have all the information VT and GSB had to label them as malicious we take a conservative stance and remove the malicious label from them. Accordingly, we label 1,344 WPNs as *known malicious*. Among them, 758 WPNs were part of 572 WPN clusters that we previously classified as ad campaigns (see Section 6.3.1). The remaining 586 WPN messages that led to malicious landing pages were not immediately found to belong to WPN ad clusters, as they formed separate small clusters. We will determine whether they are related to WPN campaigns later, in Section 6.3.3) after the meta-clustering step (explained in Section 5.3).

By using a “guilty by association” label propagation policy, as explained in Section 5.2, we label WPN ad campaigns as malicious if they include at least one *known malicious* WPN (remember that this policy is justified by the close similarity in content and landing page URL path between messages in the same cluster). This yielded 152 (out of 572) *malicious* WPN ad campaigns, which overall included 376 WPN (or more precisely their landing pages) that GSB or VT

missed to detect as malicious. After manually inspecting these 376 WPN ads, we were able to confirm that 367 of them are indeed malicious ads that lead to survey scams, phishing pages, scareware, fake alerts, social media scams, etc. We were not able to confirm the maliciousness of the remaining 9 ads (i.e., 2.4%) that led to different pages that welcome/thank the user for subscribing to the notification all hosted on the same IP address. The take away from the above discussion is that, using our WPN clustering approach, we were able to increase the number of confirmed malicious WPN ads from 758 to 1,125 (i.e., 758 plus 367), which represents an increase of about 50% as summarized in Table 4, first row.

Referring to the examples provided in Figure 4, in cluster WPN-C1, 35 out of the 40 WPNs were labeled as *known malicious* WPNs, according to VT. However, PushAdMiner labeled this entire cluster as *malicious*. After manually inspecting all 40 messages, we confirmed that the remaining 5 messages in the cluster were indeed related to the 35 malicious sweepstakes/survey scam ads.

6.3.3 Finding Suspicious Ads. So far, we have leveraged the labels provided by VT and GSB to identify malicious WPN ads, and label

WPN ad campaigns. Unfortunately, both URL blocklists suffer from significant false negatives, when it comes to detecting malicious landing pages reached from WPN ads. As an example, consider cluster WPN-C2, which PushAdMiner identifies as an ad campaign. This cluster contains 12 WPNs; none of which were labeled as *known malicious* according to VT. However, PushAdMiner flags this cluster as suspicious since it contains *duplicate ads* (see definition in Section 5.4) and via manual inspection we found that the WPN messages in this cluster display fake PayPal alerts that lead users to survey scam pages; therefore, we manually label the entire WPN-C2 cluster as malicious. This example demonstrates the gaps left by current URL blocklisting services, and how ineffective they could be if they were used to detect and block malicious ad notifications. Below we discuss how we use the meta-clustering approach explained in Section 5.4 to automatically identify and label more of such cases.

As described in Section 5.3, we apply a meta-clustering method to group WPN clusters that may relate to each other, as they share common landing page domains. To this end, we create a bipartite graph $\mathcal{G} = (W, D, E)$, here W is the set of all 8,780 WPN clusters we previously obtained, and D is the set of all 2,177 distinct landing page domains pointed to by WPN ads that we were able to record. By identifying and separating the connected components in this bipartite graph, we identify 2,046 WPN meta clusters. Of these, 224 contain a mix of WPN clusters that we previously labeled as *ad campaign* and other non-campaign WPN clusters. We then label all WPN messages contained in these 224 ad-related meta-cluster as *WPN ads*, thus increasing the number of WPN ads identified so far from 3,213 to 5,143.

Figure 5 provides two examples of meta clusters. Figure 5a contains as a node cluster WPN-C1 from Figure 4, as well as other 6 related WPN ad campaigns that are likely orchestrated by the same operators. This meta-cluster contains many *known malicious* WPN ad campaigns and WPN clusters, and therefore, we label all the WPN clusters in the meta cluster as suspicious. By manual inspection, we verified that all domains involved in this meta cluster host visually similar malicious pages (e.g., online survey scam pages).

Figure 5b shows another example of meta-cluster, which includes cluster WPN-C2 from Figure 4 as a node, along with 30 other related WPN ad clusters. In this meta-cluster, none of the WPN clusters (i.e., the nodes) were initially labeled as malicious by either VT or GSB. However, we manually inspected all landing pages pointed to by WPN messages including in the meta cluster, and we were able to confirm that these are indeed malicious, in that they display fake PayPal messages and alerts that lead users to survey scams and likely phishing-related pages.

Next, we consider all yet to be labeled WPN messages in a WPN meta cluster as suspicious if the meta cluster contains at least one *malicious* WPN cluster or if it contains *duplicate ad domains*, as defined in Section 5.4. Out of the 572 WPN ad campaigns identified earlier, we found that 255 of them contained *duplicate ad domains*. Accordingly, we were able to label a total of 287 out of 2,046 WPN meta clusters as *suspicious*. Further, we identified 166 (out of 572) additional WPN ad campaigns, which were not previously labeled malicious in the previous step, as *suspicious*. Overall, this translates into 1,479 suspicious WPN ads, as shown in Table 4. Following our manual verification process, we confirmed 1,280 (86.5%) of these ads

as malicious. The remaining 199 WPN ads were flagged by PushAdMiner because they were related to *duplicate ad domains*. Of these, 166 were alerts related to job postings and led to similar pages on multiple domains listing the same job; 23 led to multiple sites that hosted content related to the horoscope; 4 led to adult websites; and 6 were subscription welcome/thank you notifications. Notice that while these 199 WPN messages may be benign, PushAdMiner helped us identify and characterize a large number of additional WPN ads that are in fact malicious and were not identified by URL blocklists such as VT or GSB. However, our current system is not designed to be an automatic malicious WPN ad detection system. In our future work, we plan to leverage the lessons learned from the measurement results obtained in this paper to investigate how malicious WPN messages can be accurately detected and blocked in real time.

Singleton Clusters: Our tight first-stage clustering yielded 7,731 singleton clusters. Of these, 6,876 were found to share landing domains with WPNs in non-singleton clusters. After meta clustering, we were then left with 855 singleton clusters. By manually inspecting a sample of 200 singleton clusters, we found them to be a mix of simple alerts and spurious suspicious ads. Table 5 shows a few example of the text and domains related to the analyzed singleton clusters.

Table 5: Examples of Singleton Clusters

Title	Body	Domains (S)-Source (L)-Landing
TechNewsGadget	189 Fortnite Wants To Add Bot Players	technewsgadget.net (S)(L)
Congrats Walmart User!!	(1) Reward Waiting!	healthydreamstoday.com (S) besthealthlife.com (L)
The Mattest Blackest Liner EVER !!	Our new obsession!	hudabeauty.com (S)(L)
Coca Cola is looking for YOU ??	No experience required Training Provided!	eblog.network (S)(L)
Hire Local Service Professionals For All Your Needs	vconnect.com is your one stop destination for local services.	m.vconnect.com (S)(L)
FOX NEWS	Lose 45lbs In 4 Weeks! No Exercise!	nodilax.com (S) women-lifestyledaily.com (L)

Additional recent measurements: To measure the prevalence of WPN ads at a later point in time, compared to our initial measurements, we collected an additional and more recent batch of data for 5 days between April 4th, 2020 and April 9th, 2020. We revisited 300 websites randomly chosen from our previous datasets, 35 of which sent us 305 notifications over 5 days. Of these 305 notifications, PushAdMiner labeled 198 WPN ads and flagged 48 of them as malicious, which we also verified via manual analysis. After checking the corresponding landing page URLs on VirusTotal, only 15 of them were flagged as malicious, confirming again that WPN-based threats often remain undetected by current defenses.

6.4 Push Ad Networks and Blocking

Figure 6 shows the distribution of WPN ads, including malicious ones, per ad network. As it can be seen, many of the ad networks we considered in our measurements are abused to distribute malicious WPN ads.

We also investigated whether ad blocker extensions used by desktop browsers may be effective against WPN ads. For instance, we

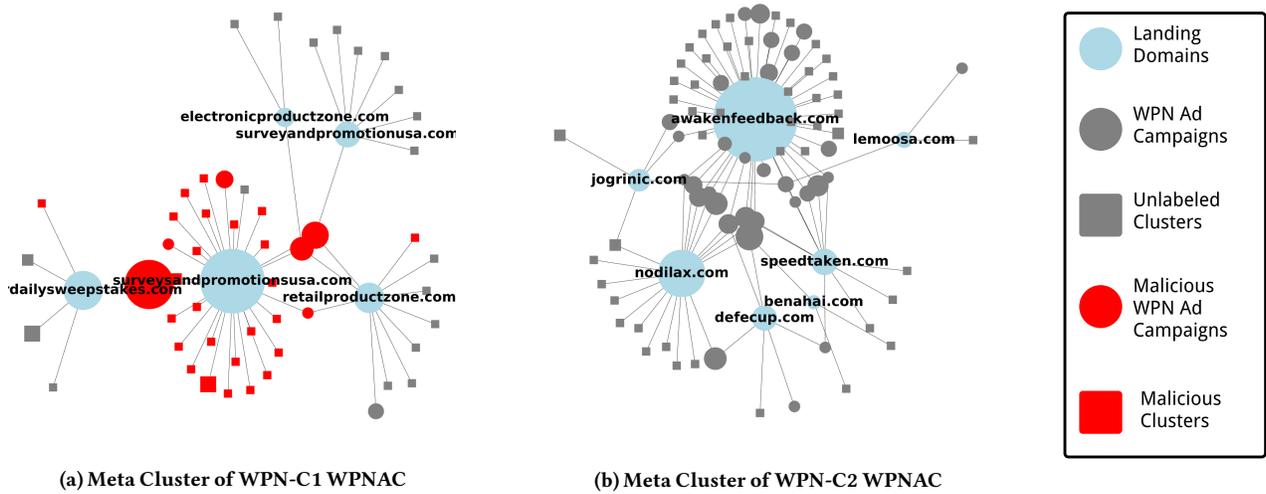


Figure 5: Graphical Representation Examples of Meta Clusters

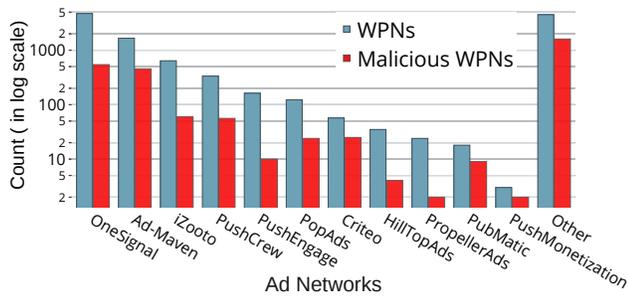


Figure 6: Distribution of WPNs w.r.t. Ad Networks

checked the URLs of SW scripts against the Easylist filter rules[6] used by most popular ad blockers. Furthermore, we installed two highly popular ad blocker extensions in our Chromium browser and checked its blocking capability. As shown in Table 6, both ad blocking mechanisms couldn't block the requests issued by the installed SW that were related to WPN ads, even though Easylist was able to filter a small number (less than 2%) of such network requests. This shows that existing methods were not sufficient to mitigate WPN-based ads, including malicious ones.

While working on this study, a new feature was introduced in Chrome [10] in February, 2020 to prevent the abuse of WPNs. This feature focuses on blocking notification permission prompts from websites that have a low notifications opt-in rate. To test the effect of this feature on our dataset, we used the latest Chrome version (Chrome 80) to revisit 300 randomly chosen websites that had previously requested notifications. Since our Chromium code instrumentations could not be easily ported to this latest version of the browser, we performed detailed manual analysis. We found

that all of the websites we visited were still able to request notification permissions without getting blocked. It is possible that this new Chrome feature will be able to block abusive WPNs in the future, after more training data is collected. However, it is unclear whether and to what extent it will be effective in blocking WPN-based malicious ads.

Table 6: Results on Existing Ad Blockers

	No. of Blocked URLs	
	Service Worker Scripts	Service Worker Requests
Easylist Blocklist	0 out of 1187	132 out of 8031
AdBlockPlus	0 out of 884	0 out of 7276
AdGuard	0 out of 895	0 out of 7520

7 RELATED WORK

In this section, we discuss prior studies related to our work.

Service Workers and WPNs: Papadopoulos et al. [48] propose multiple attack techniques to register malicious SWs that can control a victim's web browser (e.g., steal computation power for performing nefarious activities). There exist other studies [47, 51] that demonstrate similar attacks on Web Workers. Lee et al. [38] study the risks posed by PWAs. For instance, they collect SW scripts from top-ranked website homepages and analyze the risks for phishing due to the use of brand logos in push notifications. They also found a number of push-based phishing attacks that exploit WhatsApp and YouTube icons. Furthermore, they also identify security flaws in the design of browsers and third-party push notification libraries. Other works related to Progressive Web Apps[25, 42, 43] focus on the performance of PWAs in multiple environment and their energy usage.

Our study is different, because we focus on measuring the use of WPN messages as an ad-delivery platform in general, and on measuring abuse in WPN ads. To that end, we design and implement an

automated system to subscribe to, collect, interact with, and analyze in-the-wild WPN messages on both desktop and mobile platforms. **Analyzing Online Ads:** A line of studies [22, 24, 49] is dedicated to investigate the online advertisement ecosystem. They trace the information passed between ad exchanges and analyze the revenue collected using ads. Another body of studies [26, 28, 34, 40, 41, 44, 50, 55, 56, 61] focus on identifying *traditional* malicious web advertisements and ad campaigns in both desktop and mobile environments. Similarly, a recent study [58] reports observations of social engineering attack campaigns delivered through traditional web-based advertisements. Our work is different, because we specifically focus on ads and malicious ad campaigns delivered through WPN messages.

AdGraph [35] is an advanced machine learning-based system that focuses on identifying online ads embedded in web pages. AdGraph's instrumented browser collects fine-grained details about relationships between DOM nodes. A graph generated from those details is fed into a Random Forest classifier to detect ADS vs. NON-ADS elements. Upon testing the effects of AdGraph's open-source implementation on ads delivered via WPNs, we observed that their system collects detailed information from a web page's context but does not appear to record or analyze information related to SWs and WPNs. Our work is different, in that we focus specifically on collecting WPN-based ads and measuring the prevalence of malicious ads within them.

Online Scams and Ad Blockers: A number of approaches [23, 30, 59] have been proposed to study the effectiveness of existing ad block techniques as well as the various counter-measures used by ad providers to circumvent ad blockers. Other studies [27, 37, 46] discuss various online scam techniques and the prevalence of Internet fraud. Although PushAdMiner does not focus on blocking malicious ads or online scams, we demonstrate that existing ad blocking techniques are not effective at defending against malicious WPN ads.

8 DISCUSSION AND LIMITATIONS

Blocking Malicious WPN Ads: In Section 6.4, we showed that existing popular ad blockers and filtering rules were not effective at mitigating WPN ads, due to the limited visibility that extensions had on SW activities. This appears to have been due to Chromium bugs that have since been fixed (or are in the process of being fixed) [7]. However, even though extensions may now be able to observe and block SW requests, it is unclear if this will be sufficient to block WPN ads or if other evasion mechanisms can be implemented by push ad networks.

A browser extension called AdBlaster [2] specifically claims to block WPN ads. However, AdBlaster simply disables *all* push notifications, including benign non-ad ones, which can obviously disrupt legitimate uses of WPN messages. Although our PushAdMiner system is not currently designed to detect malicious WPN ads in a fully automated way, we believe that the results from our study provide useful insights into the malicious use of WPN ads that may encourage and help the security research community to build better defenses against the abuse of web push notifications.

Double Permissions: In this study, we observed some websites do not directly ask for permission to send notifications to the browser. Rather, they first create a dynamic JavaScript-based prompt that

mimics a browser permission request. Popularly referred to as *double permission*, this mechanism is used by some websites to avoid being permanently blocked from requesting notification permissions in the future (which would happen if the user clicks on the "Block" button), or to bypass other policy restrictions. We observed *double permission* to gain popularity only recently, months after our initial data collection. To check its effects on our dataset, we ran an experiment on a random sample of 200 URLs that previously requested permission directly. We found that only around 1/4th (49 out of 200 URLs that we manually checked) switched to using double permission. Furthermore, this does not impact our analysis on the abuse of WPNs. For instance, PushAdMiner could be instrumented to bypass this mechanism by simply granting notification permission to *all* visited sites, regardless of whether they request it or not. Because double permission dialog boxes are typically programmed to appear only when the website doesn't already have notification permission, they could be avoided altogether.

Evading Crawling Detection : Websites may attempt to detect browser automation to avoid being crawled. To mitigate this issue, we built our crawlers using Puppeteer. Unlike other tools, Puppeteer (developed by the Chromium team) leverages DevTools and is less prone to fingerprinting. During our experiments we also observe that a few ad networks use cookies or other information to track the device or web browser across browsing sessions. To mitigate this and increase our chances of being presented with fresh notification permission requests from as many sources as possible, we created a separate Docker container for each URL we visited.

9 CONCLUSION

In this paper, we have studied how web push notifications (WPNs) are being used to deliver ads, and measured how many of these ads are used for malicious purposes. To enable this study, we developed a system called PushAdMiner, which allowed us to automatically collect and analyze 21,541 WPN messages across thousands of different websites. Among these, our system identified 572 WPN ad campaigns, for a total of 5,143 WPN-based ads, of which 51% are malicious. We also found that traditional ad-blockers and malicious URL filters are remarkably ineffective against WPN-based malicious ads, leaving a significant abuse vector unchecked.

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for their constructive comments and suggestions on how to improve this paper, and Prof. Paul Barford for serving as our shepherd. This material is based in part upon work supported by the National Science Foundation (NSF) under grants No. 1916500 and 1909856. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

REFERENCES

- [1] 2020. Ad Block Plus. <https://adblockplus.org>. (Last accessed Sep.17, 2020).
- [2] 2020. AdBlaster Adblocker. <https://www.adblaster.org/>.
- [3] 2020. Bing Ad Content Policies. <https://about.ads.microsoft.com/en-us/resources/policies/ad-content-policies>. (Last accessed Sep.17, 2020).
- [4] 2020. Brave Ad Block. <https://brave.com>. (Last accessed Sep.17, 2020).
- [5] 2020. Cost per mille. https://en.wikipedia.org/wiki/Cost_per_mille. (Last accessed Sep.17, 2020).

- [6] 2020. Easylist. <https://easylist.to/>. (Last accessed Sep.17, 2020).
- [7] 2020. Extensions visibility into Service Worker. <https://groups.google.com/a/chromium.org/g/chromium-extensions/c/K-XAwApkyN0/m/gX9cA4ZCAgAJ?pli=1>. (Last accessed Sep.17, 2020).
- [8] 2020. Gensim Word2Vec Model. <https://radimrehurek.com/gensim/models/word2vec.html>.
- [9] 2020. Google Ads Policies. <https://support.google.com/adspolicy/answer/6020954?hl=en>. (Last accessed Sep.17, 2020).
- [10] 2020. Google Quiet UI for Notifications. <https://blog.chromium.org/2020/01/introducing-quieter-permission-ui-for.html>. (Last accessed Sep.17, 2020).
- [11] 2020. Google Safe Browsing : Blocklisting Platform. <https://safebrowsing.google.com/>.
- [12] 2020. Introduction to Push Notifications. <https://developers.google.com/web/ilt/pwa/introduction-to-push-notifications>. (Last accessed Sep.17, 2020).
- [13] 2020. Introduction to Service Worker. <https://developers.google.com/web/ilt/pwa/introduction-to-service-worker>. (Last accessed Sep.17, 2020).
- [14] 2020. Notifications API. https://developer.mozilla.org/en-US/docs/Web/API/Notifications_API. (Last accessed Sep.17, 2020).
- [15] 2020. Public WWW Usage Examples. <https://publicwww.com/examples/ads.html>. (Last accessed Sep.17, 2020).
- [16] 2020. Push API. https://developer.mozilla.org/en-US/docs/Web/API/Push_API. (Last accessed Sep.17, 2020).
- [17] 2020. Richpush Ad Network. <https://richpush.co>. (Last accessed Sep.17, 2020).
- [18] 2020. The State of Push Notification Advertising. <https://www.izooto.com/hubfs/TheStateofPushNotificationAds-iZootoreport.pdf>. (Last accessed Sep.17, 2020).
- [19] 2020. Using Application Cache. https://developer.mozilla.org/en-US/docs/Web/HTML/Using_the_application_cache. (Last accessed Sep.17, 2020).
- [20] 2020. Using Service Workers. https://developer.mozilla.org/en-US/docs/Web/API/Service_Worker_API/Using_Service_Workers. (Last accessed Sep.17, 2020).
- [21] 2020. Virus Total: Blocklisting Platform. <https://www.virustotal.com/>.
- [22] Paul Barford, Igor Canadi, Darja Krushevska, Qiang Ma, and Shan Muthukrishnan. 2014. Adscrape: Harvesting and analyzing online display ads. In *Proceedings of the 23rd international conference on World wide web*. ACM, 597–608.
- [23] Muhammad Ahmad Bashir, Sajjad Arshad, Engin Kirda, William Robertson, and Christo Wilson. 2018. How tracking companies circumvented ad blockers using websockets. In *Proceedings of the Internet Measurement Conference 2018*. ACM, 471–477.
- [24] Muhammad Ahmad Bashir, Sajjad Arshad, William Robertson, and Christo Wilson. 2016. Tracing information flows between ad exchanges using retargeted ads. In *25th {USENIX} Security Symposium ({USENIX} Security 16)*. 481–496.
- [25] Andreas Björn-Hansen, Tim A. Majchrzak, and Tor-Morten Gråvlin. 2017. Progressive Web Apps: The Possible Web-native Unifier for Mobile Development. 344–351. <https://doi.org/10.5220/0006353703440351>
- [26] Gong Chen, Wei Meng, and John Copeland. 2019. Revisiting Mobile Advertising Threats with MadLife. In *The World Wide Web Conference*. ACM, 207–217.
- [27] Jason W. Clark and Damon McCoy. 2013. There Are No Free iPads: An Analysis of Survey Scams as a Business. In *Presented as part of the 6th USENIX Workshop on Large-Scale Exploits and Emergent Threats*. USENIX, Washington, D.C. <https://www.usenix.org/conference/leet13/workshop-program/presentation/Clark>
- [28] Jonathan Crussell, Ryan Stevens, and Hao Chen. 2014. Madfraud: Investigating ad fraud in android applications. In *Proceedings of the 12th annual international conference on Mobile systems, applications, and services*. ACM, 123–134.
- [29] Docker. 2019. Docker: Enterprise Container Platform. <https://www.docker.com/>. (Last accessed Nov.1, 2019).
- [30] Kiran Garimella, Orestis Kostakis, and Michael Mathioudakis. 2017. Ad-blocking: A Study on Performance, Privacy and Counter-measures. In *Proceedings of the 2017 ACM on Web Science Conference* (Troy, New York, USA) (*WebSci '17*). 259–262.
- [31] Gensim. 2019. Documentation on Similarity Computation in Gensim Library. <https://radimrehurek.com/gensim/similarities/docsim.html> (Last accessed Nov.11, 2019).
- [32] Google. 2019. Puppeteer: Chormium Browser Automation Tool. <http://liwc.wpengine.com/compare-dictionaries/>. (Last accessed Nov.11, 2019).
- [33] Google. 2019. Set up a JavaScript Firebase Cloud Messaging client app. <https://firebase.google.com/docs/cloud-messaging/js/client>. (Last accessed Nov.1, 2019).
- [34] Chin-Tser Huang, Muhammad Nazmus Sakib, Charles Kamhoua, Kevin A Kwiat, and Laurent Njilla. 2018. A Bayesian Game Theoretic Approach for Inspecting Web-based Malvertising. *IEEE Transactions on Dependable and Secure Computing* (2018).
- [35] Umar Iqbal, Peter Snyder, Shitong Zhu, Benjamin Livshits, Zhiyun Qian, and Zubair Shafiq. 2020. Adgraph: A graph-based approach to ad and tracker blocking. In *Proc. of IEEE Symposium on Security and Privacy*.
- [36] Jordan Jueckstock and Alexandros Kapravelos. 2019. VisibleV8: In-browser Monitoring of JavaScript in the Wild. In *Proceedings of the Internet Measurement Conference* (Amsterdam, Netherlands) (*IMC '19*). 393–405.
- [37] Amin Kharraz, William K. Robertson, and Engin Kirda. 2018. Surveillance: Automatically Detecting Online Survey Scams. In *2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21-23 May 2018, San Francisco, California, USA*. 70–86.
- [38] Jiyeon Lee, Hayeon Kim, Junghwan Park, Insik Shin, and Soeul Son. 2018. Pride and Prejudice in Progressive Web Apps: Abusing Native App-like Features in Web Applications. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 1731–1746.
- [39] Bo Li, Phani Vadrevu, Kyu Hyung Lee, and Roberto Perdisci. 2018. JSgraph: Enabling Reconstruction of Web Attacks via Efficient Tracking of Live In-Browser JavaScript Executions. In *25th Annual Network and Distributed System Security Symposium, NDSS 2018, San Diego, California, USA, February 18-21, 2018*.
- [40] Zhou Li, Kehuan Zhang, Yinglian Xie, Fang Yu, and XiaoFeng Wang. 2012. Knowing your enemy: understanding and detecting malicious web advertising. In *Proceedings of the 2012 ACM conference on Computer and communications security*. ACM, 674–686.
- [41] Bin Liu, Suman Nath, Ramesh Govindan, and Jie Liu. 2014. {DECAF}: Detecting and Characterizing Ad Fraud in Mobile Apps. In *11th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 14)*. 57–70.
- [42] Tim A. Majchrzak, Andreas Björn-Hansen, and Tor-Morten Grønli. 2018. Progressive Web Apps: the Definite Approach to Cross-Platform Development?. In *HICSS*.
- [43] Ivano Malavolta, Giuseppe Procaccianti, Paul Noorland, and Petar Vukmirovic. 2017. Assessing the Impact of Service Workers on the Energy Efficiency of Progressive Web Apps. In *Proceedings of the International Conference on Mobile Software Engineering and Systems, MOBILESoft '17, Buenos Aires, Argentina, May, 2017*. to appear.
- [44] Rima Masri and Monther Aldwairi. 2017. Automated malicious advertisement detection using VirusTotal, URLVoid, and TrendMicro. In *2017 8th International Conference on Information and Communication Systems (ICICS)*. IEEE, 336–341.
- [45] Joseph Medley. 2019. Web Push Notifications: Timely, Relevant, and Precise. <https://developers.google.com/web/fundamentals/push-notifications>. (Last accessed Nov.1, 2019).
- [46] Najmeh Miramirkhani, Oleksii Starov, and Nick Nikiforakis. 2017. Dial One for Scam: A Large-Scale Analysis of Technical Support Scams. In *24th Annual Network and Distributed System Security Symposium, NDSS 2017, San Diego, California, USA, February 26 - March 1, 2017*.
- [47] Yao Pan, Jules White, and Yu Sun. 2016. Assessing the threat of web worker distributed attacks. In *2016 IEEE Conference on Communications and Network Security (CNS)*. IEEE, 306–314.
- [48] Panagiotis Papadopoulos, Panagiotis Iliia, Michalis Polychronakis, Evangelos P. Markatos, Sotiris Ioannidis, and Giorgos Vasiliadis. 2019. Master of Web Puppets: Abusing Web Browsers for Persistent and Stealthy Computation. In *26th Annual Network and Distributed System Security Symposium, NDSS 2019, San Diego, California, USA, February 24-27, 2019*.
- [49] M. Zubair Rafique, Tom van Goethem, Wouter Joosen, Christophe Huygens, and Nick Nikiforakis. 2016. It's Free for a Reason: Exploring the Ecosystem of Free Live Streaming Services. In *23rd Annual Network and Distributed System Security Symposium, NDSS 2016, San Diego, California, USA, February 21-24, 2016*.
- [50] Vaibhav Rastogi, Rui Shao, Yan Chen, Xiang Pan, Shihong Zou, and Ryan Riley. 2016. Are these Ads Safe: Detecting Hidden Attacks through the Mobile App-Web Interfaces.. In *NDSS*.
- [51] Michael Rushanan, David Russell, and Aviel D Rubin. 2016. Malloryworker: stealthy computation and covert channels using web workers. In *International Workshop on Security and Trust Management*. Springer, 196–211.
- [52] Scikit-Learn. 2019. Documentation on Silhouette Score metric to compute distance between clusters. https://scikit-learn.org/stable/modules/generated/sklearn.metrics.silhouette_score.html (Last accessed Nov.11, 2019).
- [53] Selenium. 2019. Selenium: Web Browser Automation Tool. <https://www.seleniumhq.org/>. (Last accessed Nov.11, 2019).
- [54] Grigori Sidorov, Alexander Gelbukh, Helena Gómez-Adorno, and David Pinto. 2014. Soft similarity and soft cosine measure: Similarity of features in vector space model. *Computación y Sistemas* 18, 3 (2014), 491–504.
- [55] Aditya K Sood and Richard J Enbody. 2011. Malvertising—exploiting web advertising. *Computer Fraud & Security* 2011, 4 (2011), 11–16.
- [56] Oleksii Starov, Yuchen Zhou, Xiao Zhang, Najmeh Miramirkhani, and Nick Nikiforakis. 2018. Betrayed by your dashboard: Discovering malicious campaigns via web analytics. In *Proceedings of the 2018 World Wide Web Conference*. International World Wide Web Conferences Steering Committee, 227–236.
- [57] LBE Tech. 2019. Parallel Space - Multiple accounts and Two face. <http://parallel-app.com/>. (Last accessed Nov.1, 2019).
- [58] Phani Vadrevu and Roberto Perdisci. 2019. What You See is NOT What You Get: Discovering and Tracking Social Engineering Attack Campaigns. In *Proceedings of the Internet Measurement Conference*. ACM, 308–321.
- [59] Antoine Vastel, Peter Snyder, and Benjamin Livshits. 2018. Who Filters the Filters: Understanding the Growth, Usefulness and Efficiency of Crowdsourced Ad Blocking. *CoRR* abs/1810.09160 (2018).
- [60] Wikipedia. 2019. What is Banner Blindness? https://en.wikipedia.org/wiki/Banner_blindness. (Last accessed Nov.11, 2019).
- [61] Apostolis Zarras, Alexandros Kapravelos, Gianluca Stringhini, Thorsten Holz, Christopher Kruegel, and Giovanni Vigna. 2014. The dark alleys of madison avenue: Understanding malicious advertisements. In *Proceedings of the 2014 Conference on Internet Measurement Conference*. ACM, 373–380.