Independent Auditing of Online Display Advertising Campaigns

Patricia Callejo
Universidad Carlos III de Madrid
IMDEA Networks
patricia.callejo@imdea.org

Ruben Cuevas, Angel Cuevas
Universidad Carlos III de Madrid
{rcuevas,acrumin}@it.uc3m.es

Mikko Kotila
Botlab
mailme@mikkokotila.com

ABSTRACT

The reported lack of transparency of the online advertising market may seriously affect the interests of advertisers. In this paper, we present a novel methodology that allows advertisers to independently assess the quality of display advertising campaigns. This methodology also serves to audit the accuracy and completeness of reports delivered by the vendor responsible for running a campaign. We have applied our methodology in 8 display ad campaigns configured in Google AdWords, which overall produced 160K ad impressions displayed in more than 7K publishers. Our results reveal that AdWords seems to provide incomplete information to advertisers. Specifically, we found that: (i) AdWords did not report 57% of publishers where ad impressions from our campaigns were delivered, (ii) AdWords reports a large fraction of contextually meaningful impressions based on (non-disclosed) criteria different from the publisher's theme, (iii) higher CPM investment does not lead to get impressions delivered to more popular publishers, (iv) AdWords does not offer default control of *frequency cap*, (v) around 10% ad impressions in two of our campaigns were delivered to IP's from Data Centers. The industry considers these IPs to be likely related to fraud. These findings should contribute to open a debate between advertisers and Ad Tech vendors to standardize the utilization of independent auditing methodologies as the one presented in this work.

1. INTRODUCTION

Many Ad Tech companies make the argument that online advertisements provide an effective form of advertising, and that such advertisements provide a plausible alternative to TV and other forms of traditional advertising. As a result, online advertising attracted a total investment of $125B in 2014 and it is expected to attract $240B in 2019, with an annual growth rate of 12.1% over the period [9]. At the same time, a credible body of evidence supporting the assumption on the alleged effectiveness of online advertising in comparison to other forms of advertising, is largely missing. It may very well be that other forms of media present in major advertisers' media-mix, such as TV, are far more effective than average online advertisements. The three main arguments on behalf of online advertising, that it is more accessible, lower priced per unit, and easily executable at any scale, are also the factors that have contributed to an opaque, and poorly understood fragmented market place. With thousands of vendor companies, helping advertisers place ads on millions of sites, to target over 3 billion Internet users, the online advertising ecosystem is far from transparent. Without transparency, it is not possible to truly establish if online advertising is as effective as a form of advertising as the total dollar investment in it suggests.

In particular, the opacity of this market forces advertisers to rely in reports and metrics provided by different vendors such as Ad Networks, Demand Side Platforms (DSPs) or Agency partners to assess the quality of their advertising campaigns. Some recent works have shown that, protected by this opacity, some vendors are providing inaccurate information to advertisers about their advertising campaigns [26]. These findings urge to define methodologies to allow advertisers to independently assess the quality of their online advertising campaigns as well as auditing the reports received from vendors. The research community has contributed techniques to evaluate the efficiency of different vendors in the detection and filtering of fraud [22, 26, 27, 31]. Unfortunately, fraud is not the only one aspect of the transparency problem.

In this paper, we present a lightweight and scalable methodology to audit the performance of display advertising campaigns. In essence, we propose to inject a light *JavaScript* code in our ads, a method which is typically used for collecting behavioral targeting data from a user that sees the ad. This code collects relevant information associated with each impression and sends it to a central server. Specifically, the

HotNets-XV, November 09-10, 2016, Atlanta, GA, USA

 ISBN 978-1-4503-4661-0/16/11...$15.00

DOI: http://dx.doi.org/10.1145/3005745.3005752

JavaScript code obtains the User-Agent receiving the impression, the URL where the impression was shown and user interactions with the ad impression (mouse movements or clicks on the ad). Moreover, we use the connection established with the server to obtain the IP address of the device receiving the ad impression as well as the timestamp associated to the impression. Finally, we estimate the exposure time of the ad impression as the duration of the connection.

Processing this information for an ad campaign, an advertiser would be able to objectively evaluate important quality aspects such as: (i) the potential exposition to *Brand Safety* violation episodes, (ii) the popularity and *contextual* relevance of publishers where ad impressions were delivered, (iii) the quality of delivered impressions as measured by de-facto standard metrics such as *viewability* or *frequency cap* and (iv) the exposure of the ad campaign to fraud.

We have tested the proposed methodology in 8 different campaigns set up using Google AdWords. In total these campaigns delivered around 160K ad impressions across more than 7K publishers. The obtained results indicate that the information reported by AdWords to advertisers is incomplete. In particular, our auditing methodology reveals the following insights: (i) AdWords did not report 57% of the publishers where ads from our campaigns were delivered. Without a complete list of publishers, an advertiser cannot optimise its *Brand Safety* protection; (ii) AdWords reports a large fraction of contextually relevant ad impressions based on (non-disclosed) criteria different from the publisher's thematic context; (iii) We configure campaigns with Cost-Per-Mille (CPM) investment ranging between 0,01€ and 0,30€ and conclude that, contrary to our expectation, a higher investment does not lead to impressions delivered to more popular publishers; (iv) AdWords does not impose any default *frequency cap*. This leads to hundreds of cases in our campaigns where a user receives the same ad more than 100 times with inter-arrival times between two consecutive ad impressions lower than 1 minute; (v) $\sim$10% impressions are served to IP addresses belonging to Data Centers in two of the campaigns. Note that the Ad Tech industry considers Data Center traffic to be likely associated to fraud [3, 12].

In summary, this paper contributes a novel research methodology whose application in a real use case provides solid evidences about the inconsistency of reporting from vendors in the online advertising market and how this may affect the interests of advertisers.

2. BACKGROUND

The current online advertising ecosystem is quite complex [17]. Several intermediaries (Media Agencies, Trading Desks, Demand Side Platforms -DSPs-, Ad Exchanges, Ad Networks, etc), interact in order to display ads from advertisers on publisher websites. Such intermediaries operate proprietary technologies that prevents advertisers from independently assessing the quality of advertising campaigns. Instead, advertisers have to trust the reports from their Ad Network, DSP or Agency Partner. The results in this paper suggest that the extend to which misreporting takes place may be significant.

Some of the most important aspects to consider in order to assess the quality of advertising campaigns are:

- Brand Safety: *It refers to practices and tools allowing to ensure that an ad will not appear in a context that can damage the advertiser's brand* [5]. For instance, avoiding an ad from a toy brand to be displayed on a porn website. One of the "golden rules" for an advertising campaign is to preserve the advertiser's brand safety.

- Context: Advertisers are in general interested in displaying their ads with publishers whose content is topically relevant with the topic of the ad. For instance, a hotel ad is better placed on websites related to holidays or travel agencies than on websites related to job seeking. Note that recent forms of online advertising, such as Online Behavioural Advertising (OBA) [10, 15, 20], have led to ad placements being based decreasingly in contextual relevance.

- Publishers' popularity: The popularity of a publisher indicates its capacity for attracting users. Together with other factors, it is widely used to asses the quality of a publisher. In general, advertisers pay higher CPM (cost per thousand impressions) and CPC (cost per click) for impressions placed (or clicks occurring) in popular publishers. The term *premium inventory* is generally used to describe inventory from popular websites.

- Impressions' quality: *viewability* [18] have quickly become the standard for reporting impression delivery. Based on this metric, an impression is considered to be of good quality (and thus monetized) if the user is seeing at least 50% of the pixels in the ad for at least 1 second.

Another important metric to measure the quality of an impression is the *frequency cap* [3, 11, 21, 28], which defines a limit for the number of impressions of the same ad that should be shown to the same user in a given period of time.

- Fraud indicators: World Federation of Advertisers defines advertising fraud as events "associated with an activity where impressions, clicks, actions or data events are falsely reported to criminally earn revenue, or for other purposes of deception or malice". The Interactive Advertising Bureau estimates that advertisers lose more than $8B annually directly to ad fraud in US [19].

- Conversion Ratio: The fraction of the sum of impressions that lead to a desired action (e.g. a seat booking from the airlines ticketing site).

3. METHODOLOGY

We have designed a methodology focused in HTML5 display ads, which are expected to become the de-facto standard in display advertising [7]. HTML5 allows creating ads using web technologies such as CSS or *JavaScript*. We leverage this opportunity by injecting a simple *JavaScript* code into HTML5 display ads that we buy trough an Ad Network. This code collects information about displayed ad impre-

Campaign ID	# Impressions	# Publishers	Start date	End date	CPM	Targeted Keywords	Targeted Location
Research-010	5117	350	29 March	31 March	0.10 €	Research	Spain
Research-020	42399	1777	29 March	31 March	0.20 €	Research	Spain
Football-010	33730	1086	02 April	03 April	0.10 €	Football	Spain
Football-030	24461	1367	02 April	03 April	0.30 €	Football	Spain
Russia	4096	274	29 March	31 March	0.01 €	Research	Russia
USA	1178	136	29 March	31 March	0.01 €	Research	United States
General-005	8810	580	15 February	23 February	0.05 €	Universities, Research, Telematics	Spain
General-010	42357	1549	18 February	23 February	0.10 €	Universities, Research, Telematics	Spain

Table 1: Description of the 8 AdWords campaigns used to test our auditing methodology.

ssions and sends it to a central server where it is properly stored in a database. The *JavaScript* code collects the following information: i) the URL of the webpage where the ad impression was displayed. Note that the domain part of the URL reveals the publisher; ii) the User-Agent receiving the ad impression; iii) user interactions with the ad. In particular, we collect mouse movements over the ad as well as click events. Moreover, we take advantage of the connection established between the device which received the ad impression and our server to obtain further information: iv) the IP address of the device receiving the ad, and thus, establishing the connection to our server[1]; v) the timestamp of the ad impression computed as the local UNIX time on the server at the instant of the connection establishment; iv) the *exposure time* of the ad computed as the duration of the connection measured at the server side.

We implement the described methodology employing widely used and lightweight technologies to guarantee efficiency, scalability and robustness. In particular, we use: (i) *plain JavaScript* for the code inserted in the ad; (ii) *the WebSocket protocol* [25] for transferring the information from the ad impression to the central server. Note that the information is transferred in the form of a string; (iii) *Node.js JavaScript library* [14] to parse and process the information received in the central server; (iv) *MySQL and Python* to store and process the collected datasets.

We notice that similar methodologies, using code inserted in Flash display ads, have been used in research to perform network measurements experiments [29, 30].

3.1 Limitations and Validation

The described methodology is directly applicable in ad formats that support *JavaScript* in a native manner, such as HTML5 ads. In other ad formats, such as images or video, this methodology would only work if the Ad Network allows to add a tracking pixel. Most Ad Networks and other trading platforms allow placement of 3rd-party javascript inside ads for collecting users' behavioural targeting data.

Moreover, most Ad Networks insert ads in a single (or a double) iFrame, therefore our *JavaScript* code will run inside this iFrame. There exists a widely extended security policy referred to as *Same-Origin* policy (SOP) [16], which avoids a code running as part of an iFrame tracking the activity in other parts of the webpage different from such iFrame. Hence the SOP avoids that our methodology collects information such as the upstream referrer (i.e., the website from where the user reached the current publisher). It also prevents us from collecting the position of the iFrame in the webpage, so that we cannot assess if the ad (or part of it) was shown in the visible part of the screen. This limits our methodology to measure an upper bound of the *viewability* metric presented in Section 2. This is, whether the ad was displayed more than 1 sec, but without knowing if (at least) 50% of it was shown.

We have tested our methodology in a lab controlled environment and confirmed its capacity to retrieve all the data described above. However, our methodology is expected to run in operational network environments and thus it is subject to different errors. Then, we cannot guarantee to retrieve information from every ad impression. Errors happening in the browser (e.g., untrusted *JavaScript* code not allowed to run due to the browser configuration or by an antivirus software), the network, our server, or in the connection establishment process would result in the affected ad impression(s) not being logged in our central server.

4. REAL USECASE

4.1 Ad Network and Datasets

We have applied our auditing methodology to campaigns configured in Google AdWords, which uses Google Display Network (GDN) to deliver display ads. We have selected this Ad Network due to the following two reasons: First, GDN is the most widely used Ad Network worldwide. It spans over 2 million publishers that reach over 90% of Internet users [2]; Second, GDN allows to run low budget campaigns, starting at few dollars. Then, using AdWords/GDN, we can test our methodology while respecting our budget restrictions. The main reason why we did not test our methodology in other Ad Networks is that they typically request an initial investment in the order of few thousands dollars prior to running the first campaign. This exceeded our available budget for this research.

To test our methodology we have run 8 different display advertising campaigns using Google AdWords. Overall we registered around 160K ad impressions distributed across approximately 7K publishers. We set-up campaigns with

[1]Note that we use the IP address to extract meta-data information such as the Internet Service Provider association with a user. Afterwards, we anonymize the IP using hashing techniques.

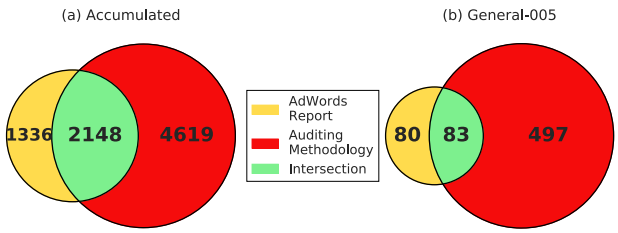

Figure 1: Venn diagram showing the number of publishers exclusively reported by our auditing methodology (red), exclusively reported by AdWords (yellow) and reported by both (green) for all our campaigns and campaign *General-005*.

different duration, different CPM values as well as different targeted keywords and geographical locations. This diversity aims at reducing the chances that observed results are due to a specific campaign set-up. Table 1 summarizes the main properties of each campaign.

4.2 Results

In this subsection, we prove the validity of our methodology to first, perform a quality assessment for our 8 display ad campaigns and, second, audit the ad campaign reports from AdWords. To this end we study the different quality aspects presented in Section 2: *Brand Safety*, *Context*, *Publishers' popularity*, *Quality of Impressions* and *Fraud Indicators*. Our campaigns were configured based on CPM. The *conversion* analysis is out of the scope of this paper, so we leave this for future work.

Note that the results presented in the rest of this section, except for the cases of *Brand Safety* and *Context*, are obtained from the analysis of the datasets resulting from our research without considering the information available in AdWords reports.

- Brand Safety: To define an efficient *Brand Safety* strategy, an advertiser must know every publisher where ad impressions are displayed in its campaigns. For each one of the 8 ad campaigns, we have compared the list of publishers where ad impressions were displayed as reported by our methodology vs. reported by AdWords. Figure 1 shows a Venn diagram representing the total number of publishers exclusively reported by AdWords (in yellow), exclusively reported by our methodology (in red) and those reported by both (in green). In particular, the figure presents results for a specific campaign (*General-005*) as well as the aggregate results across all campaigns. The aggregate results reveal that AdWords did not report 57% of the publishers where ads from our campaigns were delivered[2]. This number can increase for individual campaigns up to 75%, as in the case of *General-005*.

Part of the impressions reported by AdWords are associated with *"anonymous.google"*. These entries correspond

[2]Note that our methodology was not able to log 16.5% of the publishers.

Campaign ID	Auditing Methodology (% impressions)	AdWords Report (% impressions)
Research-010	2.50%	2.66 %
Research-020	3.75%	3.05 %
Football-010	64.12%	100 %
Football-030	46.66%	100 %
Russia	4.10%	7 %
USA	6.28%	10.73 %
General-005	4.96%	7.36 %
General-010	6.63%	56.65 %

Table 2: Fraction of impressions delivered to contextually meaningful publishers as reported by AdWords vs. our auditing methodology.

to impressions served through Google Ad Exchange to publishers or inventory partners that want to preserve their anonymity[3] [6]. Our results show that it is invalid to argue that publishers which Adwords did not report, correspond to those associated to *"anonymous.google"*. For instance, in *General-005*, AdWords registers only 425 impressions whose associated publisher is labelled as *"anonymous.google"*, however, 497 publishers identified by our methodology were not reported by AdWords. Then, even if these 425 impressions had been distributed across 425 publishers, still 72 (14.5%) publishers had not been reported by AdWords, in this specific campaign.

Therefore, *"anonymous.google"* is not the only source explaining this discrepancy. We have verified with a major Ad Tech company that this discrepancy is most likely explained by the fact that AdWords just report viewable impressions rather than all delivered impressions. Note, that this decision may have important implications for the brand safety protection of an advertiser as we argue next. An Ad Network may display an ad impression in a potentially harmful publisher for an advertiser. Whether the ad is seen or not is out of the control of the Ad Network and depends exclusively on the user's actions. If this ad is not seen by the user, then it is not reported to the advertiser. In this situation, there exists the risk that the algorithm of the Ad Network will deliver ads to that publisher again, and as a result the user may end up seeing the ad, thus leading to a brand safety violation episode. If advertisers would have access to the complete list of publishers where ads have been placed (regardless if the ad was reported to be seen or not), they could effectively identify potentially harmful sites and blacklist them. This would help prevent potential *Brand Safety* violation episodes in the future.

- Context: AdWords support guidelines indicate that campaigns configured based on *audiences* would follow a *user-targeting* strategy. Instead, campaigns configured based on *keywords*, as it is the case with our campaigns, would follow a *contextual* strategy. This means that AdWords tries to display ads within publishers whose content is related to the targeted keyword(s), and thus contextually meaningful for the campaign. In addition, AdWords may use other factors to

[3]Note that advertisers can configure their campaigns to exclude anonymous publishers [8].

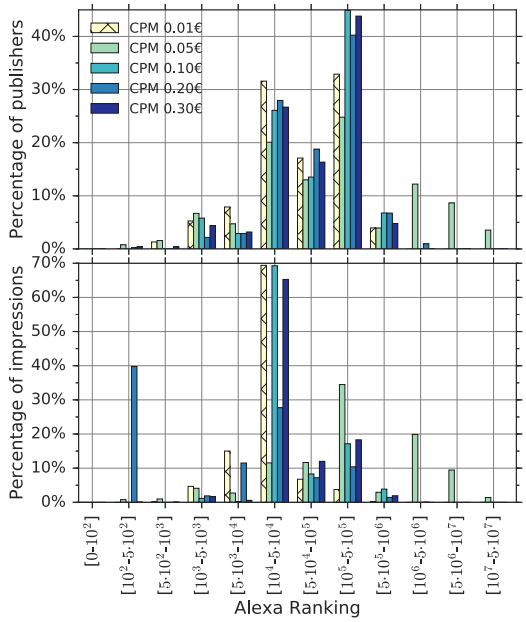

Figure 2: Distribution of publishers (top) and ad impressions (down) across the Alexa Ranking for 5 campaigns configured with different CPM investment.

determine if a publisher is contextually relevant to the campaign such as the recent browsing history of a user [1]. We have leveraged our auditing methodology to assess whether the context of a publisher is relevant to the keywords defined for a given campaign. In particular, we have obtained the keywords and topics that AdWords assigns to each publisher with at least 1 logged ad impression in our dataset. Then, we consider a publisher contextually meaningful if 1) any of its keywords match any of the campaign's keywords or 2) any of the publisher's topics are semantically similar to any of the keywords of the campaign. For this purpose we use the Leacock-Chodorow semantic similarity as described in [20].

Table 2 shows the fraction of impressions delivered to contextually meaningful publishers, as reported by AdWords vs. our auditing methodology, for our 8 campaigns. AdWords reports a notably higher fraction of ads delivered to contextually meaningful publishers compared to our methodology in most campaigns. This difference is likely due to the fact that Ad Words deliver contextual-driven impressions using other factors in addition to the publisher's theme.

- Publishers' popularity: The popularity of a publisher indicates its capacity to attract users and thus, it is one of several factors affecting the perceived quality of a publisher. In general CPMs are higher with more popular publishers, which led to our assumption that campaigns configured with a higher CPM are expected to deliver ads to more popular publishers.

Figure 2 shows the distribution of publishers and impressions across the Alexa ranking for 5 of our campaigns with CPMs ranging between 0,01€ and 0,30€. Specifically, we have defined logarithmic buckets and computed the fraction of publishers and impressions that fall in each bucket for each campaign. The results indicate that contrary to our expectation, higher CPMs do not lead to increase in impressions with popular publishers. The campaign with a CPM equal to 0,01€ seems to achieve higher than average performance with roughly 46% publishers and 89% impressions accumulated in the Alexa Top 50K sites. In comparison the campaign configured with a CPM of 0,30€ , representing a 30× investment increase, shows just 35% publishers and 68% impressions in the Alexa Top 50K. This is an unexpected observation, which may be an indication of potential inefficiencies in the market place under investigation.

Campaign ID	View $\geq 1s$
Research-010	56.18 %
Research-020	52.21 %
Football-010	79.89 %
Football-030	82.80 %
Russia	62.69 %
USA	71.13 %
General-005	75.13 %
General-010	55.03 %

Table 3: Fraction of impressions fulfilling the upper bound *viewability* criteria for each campaign.

- Quality of Impressions: In this section we evaluate the quality of impressions of our 8 campaigns using the two metrics described in Section 2, *viewability* and *frequency cap*.
Viewability: Table 3 presents the fraction of impressions that fulfills the upper bound of the *viewability* standard, and that we can measure with our methodology. The values range between 52% and 85% across campaigns. Interestingly, the two campaigns presenting the highest fraction of "viewable" impressions are the ones targeting "football", whereas other campaigns targeting other keywords (e.g., research) achieve a significantly lower viewability rate. We conjecture that the targeted context is an important factor that modulates ads *viewablity*.
Frequency Cap: Our goal in this case is to assess whether AdWords implements any default control in the *frequency cap*. Note that AdWords is used by a large number of customers without expertise in digital marketing, which may not configure a *frequency cap* in their campaigns. Therefore, it would be desirable that AdWords (or any other Ad Network) defines a default *frequency cap* on behalf of their customers. Research studies in the literature [21] have shown that a *frequency cap* over 10 does not lead to better conversion ratios. Based on this, 10 seems to be a reasonable reference value.
Figure 3 presents a scatter plot in loglog scale where the x-axis shows the number of impressions of a specific ad delivered to a user and the y-axis represents the median inter-arrival time between two consecutive impressions of that ad shown to the user. The figure presents aggregate results for all our campaigns. Note that we define a user as the combination of IP and User-Agent, so that two users behind a NAT using different User-Agents will be considered separately.
The results indicate that AdWords does not seem to use any default *frequency cap*. Indeed, 1720 (176) users receive more than 10 (100) impressions from the same ad. In addition, we

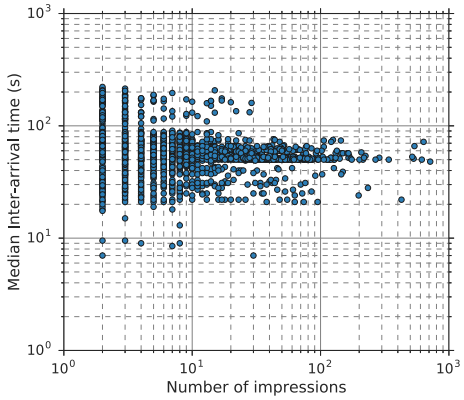

Figure 3: Number of ad impressions of a specific ad delivered to a user Vs. median inter-arrival time between impressions, considering all our campaigns.

observe that in many of these cases the inter-arrival time between impressions is rather small (below 1 min). In particular, there are extreme cases in which users receive hundreds of impressions with an inter-arrival time below 20 seconds. These observations suggest that unskilled or careless advertisers may experience inefficiencies in their campaigns performance due to the absence of a reasonable *frequency cap*.

- Fraud Identification: Fraud is one of the primary threats to effectiveness in online advertising and causes direct losses of over $8B to advertisers in US [9]. Identifying, preventing and mitigating fraud is a complex and still unsolved problem which has only recently attracted the attention of the research community [22, 23, 24, 26, 27, 31]. In this subsection we show an example of how our auditing methodology can be used to identify one common ad fraud technique. The fraud technique in question consists of installing a bot on a server. This bot can be then sent to websites to view ads or perform other revenue generating actions. Associations responsible for defining the guidelines to fight fraud such as the Media Rating Council (US) and the JICWEBS (UK) both include Data Center traffic as a common source of invalid traffic (with some exceptions such as servers that are being used for providing VPN services) and recommends vendors to filter such traffic [3, 12].

Our methodology collects the IP addresses receiving ad impressions from a given campaign. Then, we identify which of the collected IPs belong to Data Centers (e.g., Cloud Providers or Hosting Providers). We use the following methodology for this purpose: First, we used MaxMind [13] to map each IP address in our dataset to its associated provider. Second, we identified the IPs from our dataset present in a list released by Botlab [4] including more than 130M IPs belonging to the top 100 Data Center providers worldwide. Finally, for the remaining IPs, we manually verified the website of its associated provider to assess whether it offered a Data Center service or not.

Table 4 presents the results of applying the previous methodology in each of our campaigns. Specifically, it shows: (i) the fraction of IPs located in Data Centers, (ii) the portion of ad impressions delivered to those IPs and, (iii) the fraction of publishers that served impressions to those IPs. We observe that using this methodology for detection, all our campaigns deliver ad impression to Data Center IPs. Specifically, "Football" campaigns present roughly 10% of the impressions delivered to Data Center IPs and 23% of publishers exposed to such impressions. For these particular campaigns we have verified that AdWords initially charged us for more than 1K impressions delivered to Data Center IPs. Later, we got a refund from AdWords. However, AdWords did not give details on the reasons for such refund and therefore we cannot assess if the previous impressions were part of it.

Campaign ID	% of Cloud Providers IPs	% of Impressions delivered to Cloud IPs	% of Publishers showing ads to Cloud IPs
Research-010	3.39 %	4.42 %	8.62 %
Research-020	2.36 %	2.88 %	8.73%
Football-010	7.61 %	8.6 %	23.55%
Football-030	11.08 %	10.95 %	23.13%
Russia	0.52 %	0.27 %	2.58%
USA	1.03 %	0.68 %	5.56%
General-005	0.54 %	0.55 %	3.94%
General-010	0.42 %	0.58 %	2.59%

Table 4: Statistics on the volume of activity from Data Centers IPs for each campaign.

Finally, note that AdWords does not provide detailed information about the ad placement or publishers that are exposed to fraud, and thus an advertiser cannot currently assess its exposure to the analyzed type of fraud while running campaigns on Google AdWords.

5. CONCLUSION

This paper illustrates the lack of transparency and accurate information that advertisers are suffering from in the current online advertising ecosystem. This avoids advertisers from accurately assessing the efficiency and quality of their online campaigns. As a result they lack the required information to take decisions and actions to protect, for instance, their *Brand Safety*. These results should encourage advertisers to request the Ad Tech industry to standardize the use of independent measurements methodologies, as the one presented in this work. Doing so would allow advertisers to independently assess the quality of their online advertising campaigns as well as auditing the reporting practices of various vendors such as Ad Networks and DSPs.

Acknowledgments.

We would like to thank our shepherd, James Mickens, Google's team and anonymous reviewers for their valuable feedback. This work has been partially supported by the European Union through the H2020 TYPES (653449) and ReCRED (653417) projects, the Spanish Ministry of Economy and Competitiveness through the DRONEXT project (TEC2014- 54335-C4-2-R) and the Regional Government of Madrid though the BRADE project (P2013/ICE-2958).

6. REFERENCES

[1] About contextual targeting. Google Support. https://support.google.com/adwords/answer/2404186. (Date last accessed 12-October-2016).

[2] About the Google Display Network. Google Support. https://support.google.com/adwords/answer/2404190?hl=en. (Date last accessed 12-October-2016).

[3] Advantages of use frequency cap. https://www.en.advertisercommunity.com/t5/Performance-Optimization/What-s-the-advantage-of-using-frequency-capping-for-a-CPC-Ad/td-p/121528. (Date last accessed 12-October-2016).

[4] Botlab.io Deny-hosting IP List. https://github.com/botlabio/deny-hosting-IP. (Date last accessed 12-October-2016).

[5] Brand Safety Definition. http://digitalmarketing-glossary.com/What-is-Brand-safety-definition. (Date last accessed 12-October-2016).

[6] Differences between Ad Exchange and AdSense. Google Support. https://support.google.com/adxseller/answer/4599464?hl=en. (Date last accessed 12-October-2016).

[7] Display and Mobile Advertising Creative Format Guidelines. IAB, 2015. http://www.iab.com/wp-content/uploads/2015/11/IAB_Display_Mobile_Creative_Guidelines_HTML5_2015.pdf. (Date last accessed 12-October-2016).

[8] Exclude Anonymous sites. Google Support. https://support.google.com/adxbuyer/answer/159152?hl=en. (Date last accessed 12-October-2016).

[9] Global entertainment and media outlook 2015-2019. PwC, Ovum. http://www.pwc.com/gx/en/global-entertainment-media-outlook/assets/2015/internet-advertising-key-insights-1-advertising-segment.pdf. (Date last accessed 12-October-2016).

[10] IAB Europe EU Framework for Online Behavioural Advertising. IAB Europe. http://www.iabeurope.eu/files/9613/6984/1480/2012-12-11_iab_europe_oba_framework.pdf. (Date last accessed 12-October-2016).

[11] Importance of a Frequency Cap. http://www.iproacademy.com/why-you-will-fail-without-a-frequency-cap/. (Date last accessed 12-October-2016).

[12] Invalid Traffic Detection and FIltration Guidelines Addendum. http://mediaratingcouncil.org/GI063015_IVT%20Addendum%20Draft%205.0%20(Public%20Comment).pdf. (Date last accessed 12-October-2016).

[13] MaxMind GeoIP Legacy ISP Database. https://www.maxmind.com/. (Date last accessed 12-October-2016).

[14] Node.js. https://nodejs.org/. (Date last accessed 12-October-2016).

[15] Online Behavioural Advertising. Advertising Standars Authority UK. https://www.asa.org.uk/Consumers/What-we-cover/Online-behavioral-advertising.aspx. (Date last accessed 12-October-2016).

[16] Same-Origin Policy. https://developer.mozilla.org/en-US/docs/Web/Security/Same-origin_policy. (Date last accessed 12-October-2016).

[17] The Online Advertising Ecosystem Explained. Digital Ad Blog. http://digitaladblog.com/2015/02/19/online-advertising-ecosystem-explained/. (Date last accessed 12-October-2016).

[18] Viewable Ad Impression Measurement Guidelines. Media Rating Council and IAB. http://www.mediaratingcouncil.org/063014%20Viewable%20Ad%20Impression%20Guideline_Final.pdf. (Date last accessed 12-October-2016).

[19] What Is An Untrustworthy Supply Chain Costing The U.S. Digital Advertising Industry? . IAB. http://www.iab.com/insights/what-is-an-untrustworthy-supply-chain-costing-the-u-s-digital-advertising-industry/. (Date last accessed 12-October-2016).

[20] J. M. Carrascosa, J. Mikians, R. Cuevas, V. Erramilli, and N. Laoutaris. I Always Feel Like SomebodyâĂŹs Watching Me. Measuring Online Behavioural Advertising. In *Proceedings of the 11th ACM International Conference on emerging Networking EXperiments and Technologies*, CoNEXT'15, 2015.

[21] J. Chandler-Pepelnjak and Y.-B. Song. Optimal Frequency: The impact of frequency on conversion rates. Microsoft Advertising Institute, 2009. https://advertising.microsoft.com/wwdocs/user/en-us/researchlibrary/researchreport/OptimalFrequency.pdf. (Date last accessed 12-October-2016).

[22] L. Chen, Y. Zhou, and D. M. Chiu. Analysis and Detection of Fake Views in Online Video Services. *ACM Transactions on Multimedia Computing Communications and Applications (TOMM)*, 2015.

[23] V. Dave, S. Guha, and Y. Zhang. Measuring and Fingerprinting Click-spam in Ad Networks. In *Proceedings of the ACM SIGCOMM 2012 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, SIGCOMM '12, 2012.

[24] V. Dave, S. Guha, and Y. Zhang. ViceROI: Catching Click-spam in Search Ad Networks. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer Communications Security*, CCS '13, 2013.

[25] I. Fette and A. Melnikov. The websocket protocol. RFC 6455, RFC Editor, December 2011. http://www.rfc-editor.org/rfc/rfc6455.txt.

[26] M. Marciel, R. Cuevas, A. Banchs, R. González, S. Traverso, M. Ahmed, and A. Azcorra. Understanding the Detection of View Fraud in Video Content Portals. In *Proceedings of the 25th International Conference on World Wide Web*, WWW '16, 2016.

[27] B. Stone-Gross, R. Stevens, A. Zarras, R. Kemmerer, C. Kruegel, and G. Vigna. Understanding Fraudulent Activities in Online Ad Exchanges. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference*, IMC '11, 2011.

[28] Y. Yuan, F. Wang, J. Li, and R. Qin. A survey on real time bidding advertising. In *Service Operations and Logistics, and Informatics (SOLI), 2014 IEEE International Conference on*, pages 418–423. IEEE, 2014.

[29] S. Zander, L. L. Andrew, G. Armitage, G. Huston, and G. Michaelson. Investigating the IPv6 Teredo Tunnelling Capability and Performance of Internet Clients. *ACM SIGCOMM Computer Communications Review (CCR)*, 2012.

[30] S. Zander, L. L. Andrew, G. Armitage, G. Huston, and G. Michaelson. Mitigating Sampling Error when Measuring Internet Client IPv6 Capabilities. In *Proceedings of the 2012 ACM Conference on Internet Measurement Conference*, IMC '12, 2012.

[31] Q. Zhang, T. Ristenpart, S. Savage, and G. M. Voelker. Got Traffic?: An Evaluation of Click Traffic Providers. In *Proceedings of the 2011 Joint WICOW/AIRWeb Workshop on Web Quality*, WebQuality '11, 2011.