Q-Tag: A transparent solution to measure ads viewability rate in online advertising campaigns

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

Antonio Pastor
Universidad Carlos III Madrid
anpastor@it.uc3m.es

Rubén Cuevas, Ángel Cuevas
Universidad Carlos III Madrid
UC3M-Santander IBiDat
{rcuevas,acrumin}@it.uc3m.es

ABSTRACT

Viewability is one of the most important metrics used in ad-tech to measure the performance quality of ad campaigns. The viewability standard defines the visibility conditions an ad impression must meet to achieve a sufficient marketing effect to be considered viewed. The ad-tech industry offers opaque measures of viewability whose performance is questionable. To address this issue, we propose a novel methodology for measuring viewability in ad campaigns. The disclosure of the functional details of this technique makes it reproducible and auditable. Our solution has been deployed in production by a Demand Side Platform (DSP) to measure the viewability rate of the ad campaigns. Leveraging the infrastructure of this DSP, we compare the performance of our methodology with a commercial solution. Both techniques report a similar overall viewability rate of 50%. However, our solution measured the viewability in 93% of the ads served by the DSP, unlike to 74% of the ads measured by the commercial solution. A rough estimation indicates that this increase in the measured rate may lead to a revenue increase of $3.5 million per year for a mid-sized DSP serving 100M of ads per day.

CCS CONCEPTS

- **Information systems** → **Online advertising**;

KEYWORDS

Online Advertising, Quality Metrics, Viewability

ACM Reference format:
Patricia Callejo, Antonio Pastor, and Rubén Cuevas, Ángel Cuevas. 2019. *Q-Tag*: A transparent solution to measure ads viewability rate in online advertising campaigns. In *Proceedings of The 15th International Conference on emerging Networking EXperiments and Technologies, Orlando, FL, USA, December 9–12, 2019 (CoNEXT '19),* 7 pages.
https://doi.org/10.1145/3359989.3365434

CoNEXT '19, December 9–12, 2019, Orlando, FL, USA

ACM ISBN 978-1-4503-6998-5/19/12. . . $15.00
https://doi.org/10.1145/3359989.3365434

1 INTRODUCTION

The lack of transparency of the online advertising ecosystem [9, 24] has led advertisers to demand the definition of metrics that provide some guarantees in the actual marketing efficiency of served ads. To meet this demand, the ad-tech[1] industry, under the guidance of the Internet Advertising Bureau (IAB) [18] and accreditation entities such as the Media Rating Council (MRC) [8] and JICWEBS [22], has defined the viewability standard [21, 27]. Based on this standard, for instance, a display ad impression is considered *viewed* by a user only if at least 50% of the pixels of the ad are visible to the user during at least 1 second (these requirements are slightly different for other ad formats). Then, ads shown below the fold, displayed in a different tab than the one currently visible, or hidden in the background, would not be considered *viewed*. Unfortunately, as it occurs with other metrics, reported viewability rates also suffer from the opacity of the ad-tech industry. Significant stakeholders, such as Google, Facebook, or Yahoo, directly measure the viewability rate to report it to its customers. Indeed, these large vendors have defined pricing schemes that only charge their advertisers for those ad impressions meeting the viewability condition characterized by the standard [1, 2, 38]. Conversely, smaller vendors rely on third-party companies referred to as *verifiers* (Integral Ad Science [30], Moat [25], DoubleVerify [12], etc.) specialized in quality assessment of ad campaigns. All these companies use proprietary techniques to measure the viewability. As a result, the performance and limitations of such techniques are unknown. Different studies conducted by the industry and the research community have revealed episodes of inaccurate measurements of ad impressions' viewability [8, 20] as well as misreporting of different quality-related metrics [5, 10, 19]. These findings question the performance of these opaque techniques and claim for the necessity of transparent and auditable mechanisms to measure viewability (and other quality metrics).

In this paper, we propose *Q-Tag*, a novel technique for assessing if an individual ad impression meets the viewability standard criteria. Our methodology can be used to compute the viewability of individual ad impressions as well as the viewability rate of ad campaigns. The availability of *Q-Tag's* code (available under request) along with the public release of its functional details in this paper, makes it reproducible

[1] Note that along the paper we will use the term online advertising and ad-tech indistinguishably.

and auditable. We have performed a thorough evaluation of our solution through stress tests in a lab environment that report a high measurement accuracy of 93.4%.

Q-Tag has been deployed in production by a Demand Side Platform (DSP) and its performance compared in real ad campaigns with one of the most widely used viewability measurement solution in the ad-tech ecosystem. *Q-Tag* can measure viewability for 93% of the ad impressions in a campaign (on average). This represents a 19 percentage points of improvement over the commercial solution analyzed, which can measure viewability for only 74% of the ad impressions (on average). This substantial enhancement in the rate of measured ads may translate into an annual revenue increase in the order of millions of dollars for mid-size DSPs serving in the order of hundreds of millions of ads per day.

2 BACKGROUND

2.1 Overview of Online Advertising Ecosystem Operation

The online advertising ecosystem is currently responsible for delivering around a trillion ads from hundreds of thousands of advertisers into tens of millions of websites and mobile apps every day. To this end, the ecosystem has evolved into what is referred to as *programmatic advertising*. In *programmatic advertising*, the ad-spaces available in a website or mobile app are typically leased to an ad network or Supply Side Platform (SSP). The aggregated pool of ad-spaces offered by an ad network or an SSP is referred to as *ad inventory*, whereas the individual instance of an ad shown to a user is referred to as *ad impression*. Finally, the content of the ad is referred to as *creativity*. Publishers, ad networks, and SSPs form the sell side of the online advertising ecosystem since their main goal is selling ad inventory. Its counterpart, the buy-side, is formed by Demand Side Platforms (DSPs), agencies and advertisers. Advertisers typically hire the services of an agency to run their campaigns. These campaigns are configured in a DSP based on a specification including geographical location, demographic information, users' preferences, etc., from the targeted audience. Besides, the advertiser also sets up the price it is willing to pay to deliver an ad. There are two main monetization schemes: Cost Per Mille (CPM) that indicates the price an advertiser is willing to pay by 1000 impressions of its ad; Cost per Click (CPC) that indicates the price an advertiser is willing to pay if the user clicks on the ad. Finally, Ad Exchanges are the entities connecting the sell and buy sides in programmatic advertising through real-time auctions where ad-spaces from the sell-side are offered to the buy-side.

Figure 1 summarizes the described programmatic advertising ecosystem. More detailed information regarding the operation of programmatic advertising can be found in [34].

2.2 Campaign Quality Metrics

There are two main types of campaigns referred to as *branding* and *performance* campaigns, respectively. Branding campaigns aim to reach a brand or product known so that their goal is to get as many ad impressions as possible *viewed* by users. Instead, performance campaigns aim to sell a product or service, so that their goal is to persuade the user to click on the ad, bring him to the product's website, and make a purchase.

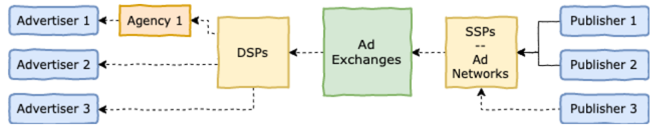

Figure 1: Overview of the programmatic online advertising ecosystem.

Since both types of campaigns have different goals, the metrics to assess their performance are also different. In branding campaigns, viewability is the key performance metric since it assesses whether the ad was sufficiently exposed to the user to have some marketing effect. In particular, the viewability standard defined by the IAB considers a display ad *viewed* if at least 50% of its pixels are exposed to the user during at least 1 second. The standard slightly differs for large display (video) ads where it is required that 30% (50%) of the pixels are shown to the user for at least 1 (2) second(s). In performance campaigns, there are two widely used metrics Return of Investment (ROI) and Click Through Ratio (CTR). Note that ROI and CTR depend on the viewability rate since the higher is the viewability rate of a campaign, the more chances to get clicks and purchases.

3 MEASURING VIEWABILITY WITH Q-TAG

Our methodology is designed to measure the viewability metric for the most common types of ads, including display and video advertisements. These ads are typically embedded in an iframe (or a nested iframe). The vendor delivering the ad controls this iframe. In addition to the ad, vendors include in the iframe the so-called *ad tags* (a.k.a. *tracking pixels*). An *ad tag* is a piece of code (typically JavaScript) that allows the vendor, or other third parties, monitoring different aspects related to an ad impression shown to a user, such as: the URL where it was displayed, the type of device receiving the ad, if there was a click event, etc. The *ad tag* sends the collected information to a server for its subsequent analysis.

We have created our JavaScript *ad tag* to measure if an ad impression meets the viewability criteria defined by the standard. We refer to it as *Q-Tag*. The straightforward manner of measuring the viewability from an *ad tag* would be to retrieve the position of the iframe in the screen and based on that, compute which fraction of the iframe is in the *viewport*, i.e., the visible part of the screen. Unfortunately, this is not (in general) possible due to a widely extended security policy referred to as the *Same-Origin Policy* (SOP) [26]. This policy would avoid our *ad tag* to retrieve the position of the iframe in the screen, in most of the cases.

To address this limitation, we have used the ability of modern browsers to stop rendering an element out of the *viewport* determined by the refresh rate, e.g., when the content is located below the fold, in a non-active tab or in the background. The refresh rate in most devices is 60 (or more) fps [14]. When an element (i.e., a pixel) is in the viewport, browsers and apps use this refresh rate. However, when the element is not in the viewport, the refresh rate pass to be close to 0, thus optimizing the use of the CPU. Hence, monitoring the refresh rate of a pixel, we can infer if it is in the viewport or not. In particular, we set up a threshold of 20 fps so that pixels refreshing at a rate equal or higher (lower) than this threshold are considered visible (not visible). We have chosen this conservative threshold to make our solution compatible in devices with overloaded CPUs that refresh at lower than 60fps rates. We have also tested our solution with thresholds of 30, 40, and 50 fps without noticing any major difference. To measure if an ad meets the viewability standard condition, we set up 25 monitoring pixels in the iframe embedding the ad and monitor the refreshing rate of each of them. The monitoring pixels are deployed in an *"X layout"* as shown in Figure 2.A: (i) ten in each diagonal (not including the central pixel), (ii) the central pixel, (iii) one pixel in each of the middle points of the four sides of the iframe ad-space (four in total). We compute the area associated with the visible monitoring pixels, and if this covers at least 50% of the area of the ad, a timer is started. If this visibility condition holds for 1 second, then we confirm that the viewability criteria has been met and the code sends an *in-view* message to the monitoring server indicating so. Contrary, if the visibility conditions change and less than 50% of the ad becomes visible before the timer reaches 1 second, an *out-of-view* event is triggered, which automatically stops the timer and restarts the process. Therefore, if the monitoring server does not receive the *in-view* message from our deployed *Q-Tag*, we conclude that the associated ad impression has not met the viewability criteria. Note that this explanation refers specifically to display ads. However, our tag can identify the type of ad (display, large display, or video) and measure the specific conditions defined by the standard for each type of ad.

4 *Q-TAG* VALIDATION

To assess the correct functionality of our solution, first, we compute the theoretical error in measuring the visible area of an ad for the selected layout and compared it with alternative ones. Second, we replicate the tests that one of the most important accreditation agencies uses to certify viewability measurement solutions. Third, we run some additional tests to analyze, among other things, the ability of our solution to measure viewability in mobile in-app ads, and in the presence of adblockers.

4.1 Layout validation

The viewability standard requires solutions that can accurately measure the viewable area of an ad and not just the viewability criteria. Based on that, the accuracy of our solution is directly associated to the selection of the number of monitoring pixels and their layout. In this subsection, we consider three different layouts: *"X layout"*, *"dice layout"*, and *"+ layout"*, whose specific deployment with 25 pixels is presented in Figures 2.A, 2.B and 2.C, respectively. Moreover, for each of these layouts, we consider deployments with a number of pixels ranging between 9 and 60. For each combination of layout and number of pixels, we compute the relative average error in the measurement of the viewable area of an ad for three scenarios: 1) *diagonal sliding*: the ad slides in the viewport diagonally; 2) *vertical sliding*: the ad enters in the screen from top to bottom; 3) *horizontal sliding*: the ad slides in the viewport from left to right. Figure 2 shows the results. If we compare the layouts, we observe that the *dice layout* offers the worst performance. The *X layout* and + *layout* offer the same performance for the vertical and horizontal sliding, but the *X layout* is the best solution in the diagonal sliding case. If we analyze now the performance as a function of the number of pixels, we observe that the error decreases fast as we move from 9 to 21 pixels, and then the error reduction flattens. The activation of a large number of pixels requires a higher computational cost without offering significant reductions in the theoretical error. 25 pixels seem to be a good trade-off offering a low error with a minimal CPU overhead.

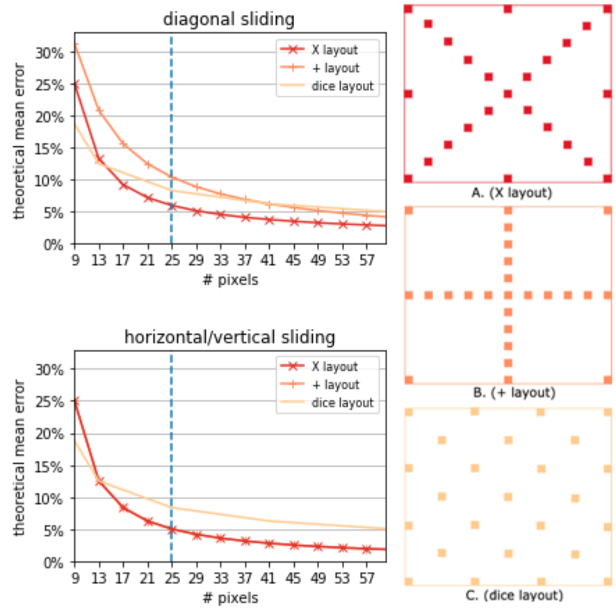

Figure 2: Comparison of possible layouts and the mean error given three scenarios for each layout.

Finally, it is worth noting that in this subsection we analyze the error in the measurement of the viewable area of an ad, which is different from measuring the viewability standard criteria. As the results in the rest of this section show, our solution offers an extremely high accuracy measuring the viewability standard.

4.2 Viewability Measurement Certification Tests

Mainly three entities define good practices in ad-tech: Media Rating Council (MRC) [8] operating in the US, JICWEBS [22] operating in the UK, and the Internet Advertising Bureau (IAB) [18] with international presence. Moreover, MRC

Test	Description	Correct result
(1) Ad within cross-domain iframes	Ad served within multiple cross-domain iframes meeting the viewability standard criteria.	The ad is always in-view and thus the solution should register an *in-view* event once the viewability criteria is met
(2) Browser is resized	The browser page is enlarged so that the ad is always *in-view* thus meeting the viewability criteria.	
(3) Out of focus	The site with the ad becomes out of focus but it is always *in-view*.	
(4) Browser moved off-screen	The browser including an ad-space is moved off-screen after meeting the viewability criteria.	The solution should register an *in-view* event once the viewability criteria is met and when the ad-space moves out of view, it should register an *out-of-view* event.
(5) Page is scrolled	The browser page including an ad-space is scrolled after the ad impression meets the viewability criteria.	
(6) Browser is obscured	The user opens another app and the ad pass to background after it meets the viewability criteria.	
(7) Tab is obscured	The user switches to a new tab within the same browser after the ad impression meets the viewability criteria.	

Table 1: Description of the tests performed by Commercial Viewability Certification

and JICWEBS developed accreditation programs to certify the correct functionality of different solutions from ad-tech stakeholders. In particular, viewability measurement solutions are subject to certification by these entities, and the list of certified providers is publicly available [31, 32]. MRC does not disclose information about its accreditation process. JICWEBS relies on a third party (ABC) to develop the viewability certification process. ABC releases its Viewability Certification report every year [39], where they describe the tests conducted for the accreditation of viewability measurement solutions. These tests analyze whether a viewability measurement solution registers the *in-view* and *out-of-view* events properly in different scenarios. Table 1 describes each one of the tests as well as the expected result from them. ABC runs these tests for two types of ads (desktop banner and desktop video) and the following pairs of the browser-Operating System: Firefox-Windows, Chrome-Windows, IE11-Windows, Safari-macOS.

Note that these certification/accreditation processes are in practice accessible only for ad-tech stakeholders, and they are expensive. Therefore, it is not feasible to obtain an official certification for our solution. Instead, we replicate the ABC tests described in Table 1 in a lab environment and confirm with ABC that our tests are indeed similar to those used in their official accreditation process. In particular, we create a testing website and an ad creativity. We embed this ad inside two cross-domain iframes[2] included in our testing website. Finally, we deploy our *ad tag* for measuring viewability within the ad creativity. Note that, we run the 7 tests used in ABC accreditation, for the same two ad formats as ABC (desktop banner and desktop video). However, we consider 6 combinations of browser-OS (two more than ABC): Firefox (v67)-Windows10, Chrome(v75)-Windows10, IE(v11)- Windows10, Safari(v12)-macOS(v10.14), Firefox68-macOS(v10.14), and Chrome(v76)-macOS(v10.14). Hence, we consider 84 different scenarios (7 test types, 2 ad formats, 6 browser-OS combinations). Note that, these pairs browser-OS represent around 82% of the current browsers market share[35]. For each of these scenarios, we automate the test process and run 500 repetitions, using Selenium WebDriver[37], except for scenarios of *test type (6)*. For these scenarios, we manually run 10 repetitions. Overall, we perform more than 36k individual tests.

The results of this thorough validation are very satisfactory since 93,4% of the 36k individual tests produce a correct result. Note that the reported 6,6% wrong results occur in *tests type (4)* and *(5)*. In those specific instances of failed tests, we are not able to register any event (*in-view* and *out-of-view*). Since this only occurs in some instances but not always, and we could not identify any consistent pattern which could explain these failures, we hypothesize the failure might be associated with the automation process with Selenium WebDriver. To check our hypothesis, we manually perform several repetitions of these tests without using the automation process, in all of them, the *in-view* and *out-of-view* events are correctly registered. Hence, we conclude that errors are more likely due to the automation process rather than the viewability measurement solution.

In summary, these results are the first reliable indication of the correct functionality of our viewability measurement solution that, in the worst case, offers a 93% accuracy.

4.3 Other tests

In this subsection, we present some extra analyses, which extend the previous validation exercise.

- *In-view event accuracy:* We randomly place a double iframe including an ad creativity embedding *Q-Tag* in 10,000 positions on the testing website. Among them, there are all sorts of cases where the ad is wholly or partially visible on the screen as well as cases in which the ad is *out-of-view*. For each one of these cases, we know the exact position of the ad on the screen and, thus, whether the *in-view* event should be triggered or not by *Q-Tag*. The results show that our solution properly triggers the *in-view* event in the 10,000 analyzed cases.
- *Mobile in-app ads:* ABC does not evaluate in-app ads in its certification process. However, based on the information publicly released by MRC, it seems it analyzes this type of ad in its accreditation process. Hence, we set up a test to evaluate that our solution correctly measures viewability for mobile in-app ads. To this end, we use the Creative Preview

[2] Note that a double cross-domain iframe is one of the most common scenarios faced by DSPs in the ad delivery process.

App from Google [33], an application for previewing mobile in-app creatives. We use this app for testing the measurement accuracy of *Q-Tag*, in the case where the ad is displayed and *in-view* in the mobile-app. We check two different creative sizes, and in both cases, *Q-Tag* notify the viewability measure correctly.

-*In-view event with adblockers and Brave:* Adblockers, as well as Brave [4], block the connection with third parties associated with ad-spaces in a webpage, and thus they block the ad delivery process. Since *Q-Tag* is only deployed if the ad is delivered, in the presence of adblockers, it should not be deployed. To confirm this, we install Adblock Plus [29] (the most popular ad blocker software) on Chrome in a lab environment and try to deliver three types of ad creativities (display, large display, or video) embedding *Q-Tag* to a testing website. We place ad-spaces in 50 random positions on the testing website for each ad type. In every test, all the connections are blocked as expected, and neither the ad nor *Q-Tag* is deployed. We reproduce the same test using Brave browser, and the ad and *Q-Tag* are not deployed, as expected.

-*Privacy-enhanced browsers:* We test our methodology in the latest Chrome, Safari, and Firefox versions (77, 13, and 69, respectively), which enable by default the prevention of cross-site tracking, i.e., blocking the third-party cookies. We reproduce the same test as in the case of AdBlock Plus and Brave. Q-Tag operates normally in these browsers since they block cookies while our methodology uses JavaScript code.

5 DEPLOYING *Q-TAG* IN PRODUCTION

Q-Tag has been deployed and integrated within Sonata, a Digital DSP/DMP Platform engineered by TAPTAP Digital[11], a multinational company with presence in more than 10 markets within Europe, North America, South America, and Africa.*Q-Tag* has been instrumented to report the viewability measures to the distributed monitoring infrastructure of this DSP. Hence, our solution is ready to be activated in any ad campaign run by this DSP. In this paper, we consider a dataset, including the viewability measures of more than 12M ads belonging to 99 ad campaigns that we monitor during a week. In addition to *Q-Tag*, the DSP allowed us to deploy the viewability measurement solution from one of the most important verifying companies in the ad-tech ecosystem[3] (also implemented as an *ad tag*). Note that the use of this verifying company has an associated cost. Due to budget limitations, we have run, both, the commercial solution and *Q-Tag*, in a subset of 4 ad campaigns including 1.89M ads.

Note that the ad campaigns considered in this paper are a representative sample of the typical operation of a stakeholder, in this case, a DSP, in the ad-tech ecosystem: 1) each of the campaigns deliver ads through several Ad Exchanges including the most important ones (AppNexus, Axonix, DoubleClick, MoPub, OpenX, Rubicon, Smaato, Smart); 2) these campaigns belong to advertisers from different sectors (e.g., *Food & Drink*, *Personal Finance*, *Style & Fashion*, etc.) and countries (e.g., US, Mexico, Colombia, Spain, UK, Germany, etc.) and thus target different audiences and geographical regions; 3) we use different size of ads (*300x250* and *320x50*) across the ad campaigns. Based on this, we believe that the performance results of our viewability measurement solution are also representative.

[3]The name of the verifying company remain anonymous to meet the terms of an NDA with the DSP.

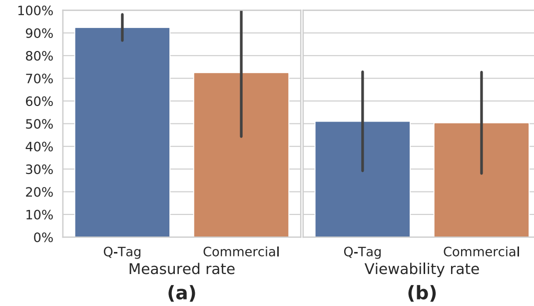

Figure 3: Comparison of the measured and viewable rate between our solution and the commercial one.

5.1 Ethics considerations

The data used in this paper has been collected by the previously referred DSP that has deployed *Q-Tag* in production. This DSP is compliant with the data protection legislation of those countries where it operates, including the recent EU data protection legislation (GDPR) [13]. Besides, the data we have received from the DSP does not include any personal information (PII) that can affect the privacy of users. Finally, the deployment of our solution is compliant with the terms of service of all providers of the DSP.

6 *Q-TAG* VS. COMMERCIAL SOLUTION

In this section, we compare the performance of *Q-Tag* and the mentioned commercial solution (one of the most widely used in the ad-tech ecosystem) using the data collected from real ad campaigns run by our DSP. In particular, we compare two performance metrics:

-*Measured rate*: This metric is defined as the fraction of ad impressions for which a solution can measure the viewability.
-*Viewability (or In-view) rate*: This metric is defined as the fraction of measured ad impressions that meet the viewability standard criteria.

Figure 3 shows the obtained results. In particular, Figure 3 (a) shows the measured rate for both solutions. The large bar shows the average, whereas the error bars show the standard deviation across the analyzed campaigns. Using this same representation, Figure 3 (b) shows the viewability rate results for both solutions.

First, we observe that both solutions offer similar average (roughly 50%) and standard deviation viewability rate. This fact indicates that our solution provides viewability rates in the same range as commercial solutions. This reinforces the conclusion regarding the high accuracy of our solution obtained through the exhaustive validation process presented in section 4.

Site type	OS	Q-Tag	Commercial Solution
App	Android	90,6%	53,4%
	iOS	97,0%	83,8%
Browser	Android	94,4%	86,7%
	iOS	94,6%	91,1%

Table 2: *Q-Tag* vs. commercial solution measured rate for site type and OS in mobile ad impressions

Second, our solution offers significantly superior performance on the measured rate. Specifically, our solution can measure (on average) the viewability for 93% ads impressions, whereas the considered commercial solution can measure just 74%. An inspection of the data reveals that most of the measurement errors of the commercial solution come from impressions delivered to mobile devices. Table 2 shows a comparison of the measured rate obtained by *Q-Tag* vs. the commercial solution sliced by the OS (Android vs. iOS) and type of site (browser vs. app). While our solution offers in any case better measured rate than the commercial one, the most significant difference occurs in the viewability measurements for Android apps, where the commercial solution can measure just 53,4% of the impressions compared to 90,6% of *Q-Tag*.

6.1 Economic implications of a higher measured rate

Based on the obtained results, DSPs can obtain an important revenue increase using our solution instead of the referred commercial one. As we mentioned above, major vendors (Google, Facebook, etc.) have opted for a pricing model that only charges advertisers for *viewed* ad impressions. The rest of stakeholders are rapidly adopting this model, so that, it is expected that shortly it will be the de-facto viewability pricing model in the ecosystem. Under this pricing model, not measured ad impressions are not monetized. In this context, a DSP using *Q-Tag* instead of the considered commercial solution would be able to measure 19% more ads. Having a 50% viewability rate reported by both solutions, roughly half of these ads would be *viewed* so that a DSP opting for our solution would effectively monetize 9.5% more ads than using the referred commercial solution. If we consider a medium-size (large) DSP serving 100M (1B) ads per day at an average CPM of $1[4], this 9.5% extra measured *viewed* ads translate into $9.5k ($95k) revenue increase per day, i.e., roughly $3.5M ($35M) per year.

7 RELATED WORK

The viewability standard was released in 2014 [27]. The wide adoption of this standard by the industry led to the development of proprietary solutions to measure viewability by verifying companies [12, 25, 30], whose performance and limitations are largely unknown. Despite the relevance of online advertising (a business generating a revenue of $107.5B in 2018 just in US [17]) and the importance of performance metrics, there is a lack of research literature addressing the viewability standard. This is probably due to the recent approval of the standard and its implementation by the ad tech industry. We could only find two theoretical studies orthogonal to our work. Chong Wang et al. have created a model to predict the viewability analyzing scroll depth for a given user and a page [36]. In a different work, David Bounie et al. [3] presented an analysis of the economic consequences of the investment in campaigns with low viewability rates.

From a measurement methodology perspective, we find previous works in the literature performing measurements from code embedded in ads. Some of these works use flash ads as a platform for measuring network properties and security aspects [28, 40, 41]. Note that most DSPs no longer support flash because it is deprecated in online advertising. More recent measurement works use JavaScript-based ad measurements for auditing the online advertising ecosystem [5, 15], for measuring mobile devices network performance [6, 7], or for measuring DNS aspects [16, 23], among others. Note that none of these works present a methodology able to measure viewability as we do in this paper.

8 CONCLUSION

In this paper, we have described, implemented, and evaluated *Q-Tag*, a new technique for measuring the viewability rate of online advertising campaigns, which offers a 93.4% measurement accuracy.

The release of functional details of our technique for measuring viewability makes it easily replicable by advertisers, agencies, or DSPs. In consequence, these stakeholders have *for first time* at their disposal an independent and auditable solution for assessing the viewability rate of their campaigns, without the need to rely upon opaque solutions offered nowadays by the industry.

Q-Tag has been deployed in production in a DSP. Using information from 12M measured ads served by this DSP, we compared the performance of our technique with one of the most important commercial solutions for viewability measurement. The comparison results show that *Q-Tag* can measure the viewability in 19% more ads than the commercial solution. A ballpark estimation reveals that these extra measured ads may lead to an annual revenue increase in the order of millions (tens of millions) of dollars for mid (large) size DSPs.

ACKNOWLEDGMENTS

We would like to thank our shepherd, Marco Mellia, and our anonymous reviewers for their valuable feedback. This work has been partially supported by Ministerio de Economía y Empresa (Spain) under project TEC2016-76795-C6-3-R and grant RyC-2015-17732, and by the European H2020 project SMOOTH (786741).

[4]Note that a $1 average CPM is a realistic reference in the ad-tech ecosystem.

REFERENCES

[1] Martin Beck. Adage. 2015. Google Display Network Will Charge Only For Ads That Are Viewed. (2015). Retrieved October 16, 2019 from http://marketingland.com/google-display-network-will-charge-only-for-ads-that-are-viewed-144768

[2] Tim Peterson. Adage. 2014. Yahoo to Charge Some Advertisers Only for Ads People Actually See. (25 June 2014). Retrieved October 16, 2019 from http://adage.com/article/digital/yahoo-charge-advertisers-ads-people/293857

[3] David Bounie, Valérie Morrisson, and Martin Quinn. 2016. Do You See What I See? Ad Viewability and the Economics of Online Advertising. (2016).

[4] Brave. 2019. Secure, Fast & Private Browser. (2019). Retrieved October 16, 2019 from https://brave.com/

[5] Patricia Callejo, Ruben Cuevas, Angel Cuevas, and Mikko Kotila. 2016. Independent Auditing of Online Display Advertising Campaigns. In *Proceedings of the 15th ACM Workshop on Hot Topics in Networks*. 120–126.

[6] Mark D Corner and Brian N Levine. 2018. MicroMobile: Leveraging mobile advertising for large-scale experimentation. In *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, 310–322.

[7] Mark D Corner, Brian N Levine, Omar Ismail, and Angela Upreti. 2017. Advertising-based Measurement: A Platform of 7 Billion Mobile Devices. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*. ACM, 435–447.

[8] Media Rating Council. 2018. (2018). Retrieved October 16, 2019 from http://mediaratingcouncil.org/

[9] Brian Dalessandro, Claudia Perlich, Ori Stitelman, and Foster Provost. 2012. Causally motivated attribution for online advertising. In *Proceedings of the Sixth International Workshop on Data Mining for Online Advertising and Internet Economy*. ACM, 7.

[10] Shailin Dhar. 2016. Mystery Shopping Inside the Ad Fraud Verification Bubble. (2016). Retrieved October 16, 2019 from http://www.slideshare.net/ShailinDhar/mystery-shopping-inside-the-adverification-bubble

[11] TAPTAP Digital. 2019. Location Intelligence for Marketing. (2019). Retrieved October 16, 2019 from http://www.taptapdigital.com

[12] DoubleVerify. 2017. (2017). Retrieved October 16, 2019 from http://www.doubleverify.com

[13] EUGDPR. 2016. The EU General Data Protection Regulation. (2016). Retrieved October 16, 2019 from http://www.eugdpr.org/

[14] Google. 2019. Rendering Performance. (2019). Retrieved October 16, 2019 from https://developers.google.com/web/fundamentals/performance/rendering/

[15] Saikat Guha, Bin Cheng, and Paul Francis. 2010. Challenges in measuring online advertising systems. In *Proceedings of the 10th ACM on Internet Measurement Conference*. 81–87.

[16] G. Huston. 2013. Measuring Google's Public DNS. (2013). Retrieved October 16, 2019 from https://labs.ripe.net/Members/gih/measuring-googles-public-dns

[17] IAB. 2019. Internet advertising revenue report. 2018 full year results. (2019). Retrieved October 16, 2019 from https://www.iab.com/wp-content/uploads/2019/05/Full-Year-2018-IAB-Internet-Advertising-Revenue-Report.pdf

[18] Internet Advertising Bureau. IAB. 2017. (2017). Retrieved October 16, 2019 from https://www.iab.com

[19] Lara O'Reilly. Business Insider. 2016. Facebook over-inflating its video view count is bad, but not as bad as it seems. (23 Sept. 2016). Retrieved October 16, 2019 from http://nordic.businessinsider.com/facebook-exaggerating-view-count-didnt-affect-advertiser-bills-2016-9

[20] Lara O'Reilly. Business Insider. 2016. Two of Google's metrics have been suspended from a key accreditation service used to measure ads. (12 Oct. 2016). Retrieved October 16, 2019 from http://www.businessinsider.com/google-doubleclick-suspended-from-media-rating-council-accreditation-2016-10

[21] JICWEBS. 2018. Viewability Product Principles. (06 July 2018). Retrieved October 16, 2019 from https://jicwebs.org/wp-content/uploads/2018/07/JICWEBS_Viewability_Product_Principles_July_2018.pdf

[22] JICWEBS. 2019. (2019). Retrieved October 16, 2019 from https://jicwebs.org/

[23] Wilson Lian, Eric Rescorla, Hovav Shacham, and Stefan Savage. 2013. Measuring the Practical Impact of DNSSEC Deployment. In *Presented as part of the 22nd USENIX Security Symposium*. USENIX, Washington, D.C., 573–588.

[24] Miriam Marciel, Rubén Cuevas, Albert Banchs, Roberto González, Stefano Traverso, Mohamed Ahmed, and Arturo Azcorra. 2016. Understanding the Detection of View Fraud in Video Content Portals. In *Proceedings of the 25th International Conference on World Wide Web*.

[25] Moat. 2017. Moat Analytics. (2017). Retrieved October 16, 2019 from https://moat.com/analytics

[26] Mozilla. 2016. Same-OriginPolicy. (2016). Retrieved October 16, 2019 from https://developer.mozilla.org/en-US/docs/Web/Security/Same-origin_policy

[27] MRC and IAB. 2014. MRC Viewable Ad Impression Measurement Guidelines. (30 June 2014). Retrieved October 16, 2019 from https://www.iab.com/wp-content/uploads/2015/06/MRC-Viewable-Ad-Impression-Measurement-Guideline.pdf

[28] Mark O'Neill, Scott Ruoti, Kent Seamons, and Daniel Zappala. 2016. TLS proxies: Friend or foe?. In *Proceedings of the 2016 ACM on Internet Measurement Conference*. 551–557.

[29] Adblock Plus. 2019. (2019). Retrieved October 16, 2019 from https://adblockplus.org/

[30] Integral Ad Science. 2017. (2017). Retrieved October 16, 2019 from https://integralads.com

[31] MRC Accredited Services and Services Under Review. 2019. (2019). Retrieved October 16, 2019 from http://www.mediaratingcouncil.org/Accredited%20Services.htm

[32] JICWEBS Signatories. 2019. (2019). Retrieved October 16, 2019 from https://jicwebs.org/certification-process/signatories/

[33] Google Support. 2019. Preview ads on a phone or tablet. Creative Preview app. (2019). Retrieved October 16, 2019 from https://support.google.com/richmedia/answer/2879163

[34] IAB UK. 2018. Back to Basics Guide to Programmatic. (2018). Retrieved October 16, 2019 from https://www.iabuk.com/standards-guidelines/back-basics-guide-programmatic

[35] W3Counter. 2019. Browser & Platform Market Share. (2019). Retrieved October 16, 2019 from https://www.w3counter.com/globalstats.php

[36] Chong Wang, Achir Kalra, Cristian Borcea, and Yi Chen. 2015. Viewability prediction for online display ads. (2015), 413–422 pages.

[37] Selenium WebDriver. 2019. Browser Automation. (2019). Retrieved October 16, 2019 from https://www.seleniumhq.org/projects/webdriver/

[38] Devay Alba. Wired. 2015. Facebook says it will not charge for an ad unless someone sees it. (2015). Retrieved October 16, 2019 from https://www.wired.com/2015/02/facebook-an-ad-doesnt-count-unless-someone-sees-it

[39] ABC with JICWEBS. 2018. Viewability Certification. (2018). Retrieved October 16, 2019 from https://www.abc.org.uk/images/Viewability_Report.pdf

[40] Sebastian Zander, Lachlan L.H. Andrew, Grenville Armitage, Geoff Huston, and George Michaelson. 2012. Investigating the IPv6 Teredo Tunnelling Capability and Performance of Internet Clients. *ACM SIGCOMM Computer Communications Review* (2012).

[41] Sebastian Zander, Lachlan L.H. Andrew, Grenville Armitage, Geoff Huston, and George Michaelson. 2012. Mitigating Sampling Error when Measuring Internet Client IPv6 Capabilities. In *Proceedings of the 2012 ACM on Internet Measurement Conference*.