Censorship Addendum for Advertising-based Measurement: A Platform of 7 Billion Mobile Devices

Mark D. Corner, Brian N. Levine, Omar Ismail, and Angela Upreti
College of Information and Computer Sciences University of Massachusetts Amherst, MA, USA
mcorner,levine,oismail,aupreti@cs.umass.edu
July 24, 2017

APPENDIX

Given the controversial nature of studying censorship, we had to remove the following information from our submission to MobiCom [4]. We received negative feedback from the MobiSys program committee on this issue and felt that it was not worth the risk in having it rejected at MobiCom. We disagree with the conclusion reached by the MobiSys TPC and in particular we do not believe that the non-interrogative nature of TPC is the best place to evaluate human subjects concerns. You are welcome to cite this work separately from the MobiCom paper.

Before conducting the research described here, we sought and obtained the consent of our own Institutional Review Board (IRB) under protocol numbers 2016-3112 and 2016-3141. Before starting the censorship study we sought additional IRB approval, as similar measurements taken from a website have been considered an area of ethical concern by some in our community [5]. In our study of censorship, users unwittingly retrieve blocked images. Burnett et al[3] used a similar technique of background image loading to measure censorship, but relied on using their own webpage. Their project was submitted to the Georgia Tech and Princeton IRBs, which declined to review the project because it was not collecting PII. Our IRB reviewed our project and has also allowed the study, with restrictions that include using what the Encore paper calls "ubiquitous, yet uninteresting URLs"[3]. These URLs are often found embedded in webpages and constitute non-controversial content (e.g., the Google logo), even if they come from sites generally blocked by governments, such as Facebook and Google; see Fig. 1.

Figure 1: The 88×88 image we retrieved from youtube, blocked in China and Iran.

A. ADVERSARIAL ENVIRONMENTS AND CENSORSHIP

Challenge: Web retrievals made by an advertisement may be blocked or severely hampered by the network environment and browser. For instance, a firewall operated by a company or a government may block or tamper with images or data requested by the browser. In addition, page access is ruled by the *same-origin policy*, which is a rule applied by all browsers disallowing a script coming from one domain from requesting data from other domains.

One of the rare exceptions to same-origin is that HTML image tags for arbitrary domains can be placed in a document, though the script cannot access the contents of those images. By placing an image tag pointing to another domain and receiving javascript callbacks on image loads, AaaP can measure how long an image takes to load, if the image was fetched correctly, and get the original (i.e., *natural*) image width and height. Using image tags, AaaP can measure if particular domains are blocked in the device's network environment, in particular Internet censorship imposed by governments.

A similar technique for embedding images was described in Encore [3] using web pages rather than advertisements. The advantages in using AaaP to measure censorship is that we have extremely fine-grained control over how and when the measurement occurs. In a web page, one has to entice users to visit the page and thus measurements may be poorly distributed and irregularly timed, whereas with advertisements specific geographic areas can be targeted to give fine grained resolution. For example, we can find if specific areas of a country are censored, such as blocks often placed in certain regions in Turkey [1].

To measure censorship AaaP includes javascript to add an image tag with a source taken from a set of image URLs we hypothesize to be censored or not censored in various countries. Based on our IRB proto-

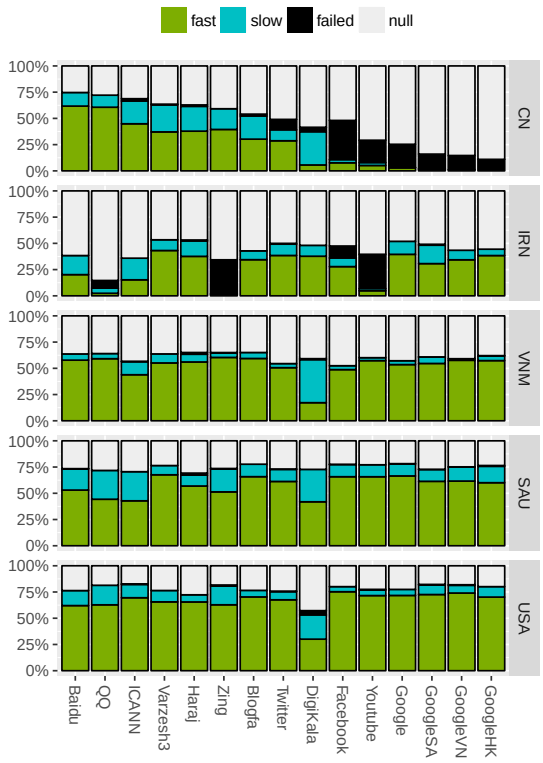

Figure 2: Measurements are marked slow if loading requires more than 3 seconds. NULL measurements means the image failed to load by the time the ad closed by the user or application. Total cost = $10.06; Impressions = 27,760; Efficiency = 100%.

col, we used only "uninteresting" and non-controversial images such as logos and user-interface elements (e.g., Figure 1). We then determine which of four results occurred: the image loaded *fast* and correctly (we verified the native height and width of the image and fetch using SSL to further verify the client received the correct image); the image loaded correctly but *slowly*; the image *failed* to load; or the image failed to load by the time the advertisement closed, termed a *null* result. The experiment comprises more than 27k measurements, costing $10. We further verified the results through testing using open web proxies in the countries. The results of the experiment, including the sites the images came from and the countries we tested are shown in Figure 2.

We can conclude a number of things from the results: *(i)* The primary result is that we can effectively measure censorship of Facebook, Google, and YouTube in China and the censorship of Youtube in Iran, as well as QQ and Zing.vn—both of these companies provide messaging and chat applications which are typical targets for censorship. In China, images typically fail to load in a timely manner, rather than producing an immediate error. Interestingly, the results show that an image from Twitter is not blocked, even though the content of Twitter is known to be blocked in China [2]. *(ii)* Images from blocked sites in China and Iran are not blocked by all ISPs. For example, Facebook was almost always allowed by ISPs in Guangdong (which contains Hong Kong) and the China Education and Research Network. *(iii)* DigiKala, a Persian shopping site exhibits poor performance, except in Iran, and the lack of access is not due to censorship, but rather an issue with poor hosting performance. Overall, AaaP is an effective, and inexpensive method for inexpensive real-time measurement of blocked and excessively hampered networks.

Most importantly, these results demonstrate a generally useful feature in AaaP: the ability to test both from and to remote resources on the Internet. There are even more techniques that we can use in AaaP, but require additional exploration, such as using CORS-enabled resources, JSON-P to test content APIs, rather than just images, and WebRTC to measure P2P accesses.

B. REFERENCES

[1] `https://turkeyblocks.org/2016/10/27/new-internet-shutdown-turkey-southeast-offline-diyarbakir-unrest/`.

[2] David Bamman, Brendan O'Connor, and Noah Smith. Censorship and deletion practices in chinese social media. *First Monday*, 17(3–5), 2012.

[3] Sam Burnett and Nick Feamster. Encore: Lightweight measurement of web censorship with cross-origin requests. *ACM SIGCOMM Computer Communication Review*, 45(4):653–667, 2015.

[4] Mark D. Corner, Brian N. Levine, Omar Ismail, and Angela Upreti. Advertising-based measurement: A platform of 7 billion mobile devices. In *Proceedings of the 23rd ACM International Conference on Mobile Computing and Networking (MobiCom)*. Snowbird, Utah, USA, October 2017.

[5] Ben Jones, Roya Ensafi, Nick Feamster, Vern Paxson, and Nick Weaver. Ethical concerns for censorship measurement. In *Proc. ACM SIGCOMM Workshop on Ethics in Networked Systems Research*, pages 17–19. ACM, 2015.