Encore: Lightweight Measurement of Web Censorship with Cross-Origin Requests – Public Review

John W. Byers
Dept. of Computer Science, Boston University
byers@cs.bu.edu

Piggybacking on web requests has long been a feature of the open Internet: witness the lengthy cookie-drop chains and calls to ad servers hanging like rows of barnacles off of most major websites. To date, most non-malicious occurrences of forking or redirecting web requests have had a notional connection to the user's intent. The Encore paper, on the other hand, proposes a novel use of piggybacking toward a worthy goal, but one that is entirely unrelated to the underlying web requests: monitoring global censorship. The key technical idea behind Encore is: to measure whether users in a given region like Switzerland can access URL X, Encore (deployed, say, at the Princeton web server) awaits an HTTP request from a Swiss user and provides javascript that contains an embedded cross-origin request for URL X as part of the HTTP reply. The Swiss user's browser then executes the request for URL X, and the conditional javascript execution allows the Princeton web server to subsequently infer whether or not the request was censored. By enlisting requests and compiling observations from a worldwide set of vantage points, Encore can monitor censorship globally with no remotely deployed infrastructure or software! The essential prerequisite is installation on an uncensored web server that has global reach. In a prototype demonstration directing remote users to download URLs from a mock censorship testbed, the authors give compelling evidence of the feasibility of this approach at scale.

The Encore paper and the approach gave reviewers pause on several levels. First was the shock factor associated with the power (and relative untraceability back to the true source) of cross-origin requests, used in this light. Monitoring global censorship is just one of many uses of injecting remote measurements that one could easily imagine, some beneficial, e.g., monitoring reachability of end-to-end paths in an enterprise network; others, such as enabling a new DDoS attack vector, not so much. But more problematic were the reviewers' ethical concerns: for Encore to work effectively at scale, unsuspecting users around the world must be enlisted to download oft-censored URLs without their informed consent. These requests could potentially result in severe harm; for example, when the user lives in a regime where due process for those seen as requesting censored content may not exist. Reviewers viewed such measurements as unethical and disavowed this Encore use case. While this ethical issue could be sidestepped by obtaining informed consent, PC members and survey respondents of an independent study [41] agreed that most users for whom censorship is an issue would be unlikely to consent to Encore's measurements.

Had the paper stopped here, I suspect this would have been a less controversial decision for the PC: fascinating idea, really quite eye-opening that this is possible, well-conceived and prototyped, and conducted in a research area (network measurement of repressive state-level activities) that the PC affirmed as highly important. But, the Encore authors touched a third rail of sorts by going beyond prototyping – they also conducted measurements in the wild. Initially, they made some measurements using censored URLs from the Herdict lists, but after consultation with ethics researchers, limited their measurements to favicons of popular (but still sometimes censored) websites, which continue to this day. The PC wrangled over the question of whether the initial measurements, viewed as unethical, should electrocute the paper, despite its technical merit. In lengthy reflections, many mitigating circumstances were laid out: the authors acted in good faith, they disclosed the existence of experiments with problematic URLs but did not rely upon or report those results in the paper, they consulted their cognizant IRBs regarding their experiments, and the SIG does not yet have formal guidelines in this space that the authors violated. So rather than silently reject the paper, the PC voted to publish, with their ethical concerns documented in the unprecedented "signing statement" atop the paper.

In the end, novel network measurement research that seeks to characterize the extent of highly sensitive government activities such as censorship and surveillance can be of great value. At the same time, the risk and potential for personal harm associated with these measurements raise new challenges and ethical issues for our community that will warrant our close attention.