Challenges in Measuring and Evading Nation-state Censors

Dave Levin
University of Maryland

Censorship by nation-states is a common and pernicious threat to free and open communication on the Internet. Multiple countries—such as China, India, Iran, and Kazakhstan—aggressively filter traffic crossing their borders, using it to limit access to information and quell political dissidents.

Considerable and impressive measurement efforts have been made towards *understanding* how these censors operate. This includes groups like OONI, Ioda, CensoredPlanet, Geneva, and others. Core to all of these efforts is measurement: it can lead to better understanding of the impact that censors have (even to those outside their borders [1]) and it can lead to better evasion strategies [2].

Although measuring censors is very important, it is also extremely challenging to do in a safe, ethical, and reliable manner. Currently, the bar is exceedingly high to get started in censorship-related research, often requiring personal connections. To be honest, I am not sure what the answer is, but my hope is that this can help drive a conversation towards a more systematic and communal approach.

How anti-censorship research is performed

One of the primary research goals is to develop a better understanding of how censors operate: what, whom, and how do they censor? Performing such measurements typically involves some way of actively issuing requests for censored content (censored domain names or keywords), and evaluating whether or not it was blocked or dropped. Various measurement efforts differ primarily in terms of what kinds of queries they send, from what vantage points, and how they detect whether censorship took place.

Additionally, there is extensive work on developing tools or protocols for bypassing censors or avoiding them altogether. Even in this space, there is a large amount of measurement work: after all, testing whether or not a new tool works often requires running against live censors. Likewise, finding new evasion strategies often sheds light on how censors (do not) work. Evasion research faces many of the same concerns as pure measurement, and amplifies some of the risks (anecdotally, censors seem not to be as offended by being measured as they are by being actively circumvented).

Some of the challenges this raises

- **Ethics:** It can be difficult to perform such measurements ethically; What vantage points are safe to use, and who faces what risks? The community has even had some trouble reaching consensus on the potential harm of certain measurement methods [3].
- **Safety of participants:** Some participants within censoring regimes graciously offer up their own computers or rent them on researchers' behalf, but this can potentially put them in danger, as well. Researchers who work in this field take extra precautions to ensure their safety, by securing their online communication and obfuscating details about their experiments.
- **Safety of student researchers:** Merely working in this area can put researchers at risk. This is of special concern for students; they must be informed of the risks before engaging in the research. The community should develop shared knowledge on what the risks are, how to communicate them to researchers (but especially students), and measurement methodologies to mitigate risks.
- **Availability of vantage points:** Some countries introduce severe restrictions on who is allowed to rent machines from cloud providers[1]. This makes it difficult to run experiments with broad coverage of a country (censoring practices can differ between ISPs).
- **Reproducibility:** The difficulty in getting access points also makes it difficult for researchers to reproduce others' results—and raises the bar of entry for the field as a whole. Shared infrastructure could help.

I am not the first to note these challenges—rather, given the nature of the work, the community has, by and large, been proactive in identifying and mitigating risks to users. However, there remain significant hurdles to being able to perform this research; I suspect that the NSF could play a role in bringing together various leaders in this field to establish a set of best practices and shared infrastructure.

References

[1] Anonymous. The Collateral Damage of Internet Censorship by DNS Injection. In *ACM SIGCOMM CCR*, 2012.

[2] Kevin Bock, George Hughey, Xiao Qiang, and Dave Levin. Geneva: Evolving Censorship Evasion Strategies. In *ACM CCS*, 2019.

[3] Sam Burnett and Nick Feamster. Encore: Lightweight Measurement of Web Censorship with Cross-Origin Requests. In *ACM SIGCOMM*, 2015.

[1] For instance, activating an account on a Chinese cloud provider requires a Chinese cell phone number, which in turn requires a Chinese bank account, which in turn requires a Chinese national ID card.