Network security, anonymity and censorship
SaTC 2019 PI meeting breakout group report
Co-leads: Nick Feamster, Nick Hopper

1. Problem/Domain Summary

Protecting metadata - concealing who communicates with whom, in order to provide free and open access to the Internet, and enable controlled identity disclosure, is an important tool to provide fundamental digital rights such as freedom of expression and freedom of association. However, such protection may be at odds with several emergent Internet and technological trends, such as:

- Increasing centralization of the Internet: having a small set of providers for important services such as DNS and CDN content means that some entities may be in a position to collect large amounts of metadata.
- Trends such as IoT and blockchain also centralize collection of private data such as financial transactions, video and audio recordings, appliance use, and so on.
- Increased use of carrier-grade NAT alongside increasing penetration of IPv6 invalidate some assumptions about 1:1 identity:address relationships

The charge of this group was to think about the challenges involved with these changes and their interplay with existing problems in the space of network anonymity and censorship.

2. Key Research Challenges

Centralization - centralized services such as DNS, content delivery networks, browser vendors, and OS platforms have enabled relatively small groups to bring about changes in content security ("encrypting the internet") but this has brought new challenges since large amounts of internet traffic are now visible to a few entities, who may also have access to client-side data as well. How does this affect the privacy of metadata and the ability to censor or selectively alter content? Its resiliency? Is it problematic that a small set of entities can "turn off" the internet? Or perhaps circumvention systems such as Tor, since many relays and VPS proxies are hosted on cloud providers? How can we evaluate the impact of this centralization and detect efforts to censor or monitor traffic by such entities?

Information Asymmetry - censors and surveillance apparatus can monitor the networks they control. This allows them to understand "normal" patterns of behavior, traffic volumes, and daily usage patterns. All of these can impact the effectiveness of a circumvention or anonymity system. How can researchers responsibly and safely obtain data sets that allow evaluation of these systems while *not* engaging in surveillance? How can we privately share these data sets? Can these data sets assist in the evaluation of new systems for anonymity and circumvention?

Modeling adversary and response capabilities - how do we predict censor response to circumvention tools? Recent events show that some censors may be willing to go to extreme measures in response to some techniques, such as domain fronting. While there is work in modeling the state of devices such as firewalls and middleboxes, understanding the incentives and response capabilities of potential censorship and surveillance infrastructures seems more difficult; and may shift which entities should be seen as trustworthy.

Measuring and observing censorship - many studies have emerged on filtering techniques such as DNS hijacking, IP blocking, DPI and active probing. Understanding other approaches to censorship such as chilling effects and surveillance is challenging since the actions are not directly observable (and may be counterfactual). Measuring "flooding"-based attacks that censor information through disinformation can be challenging as well, due to the possibility of honest disagreement between sources.

Other trends such as IoT, blockchain, content deanonymization, content moderation, carrier- grade NAT, IP6, and Internet Islands may represent both challenges and opportunities in this space as well.

3. Potential Approaches

The group discussed some potential approaches, but no clear consensus emerged:

- Measurement studies can help us understand the extent of centralization, but may not answer how these trends should change adversarial focus. Companies sometimes want to share data but need privacy assurances as well.
- NSF-Funded data sets for the purpose of sharing and evaluation are needed as a critical infrastructure for progress in this domain.
- Differential privacy may help with data sharing but generative models may not capture the "right" properties of traffic needed in later studies. Deanonymization is a substantial risk in case of some relaxed models of differential privacy.
- Trusted enclaves could assist in the generation of private data sets, and protection against adversaries and censors that control end-user or cloud platforms.
- Economics and game theory may help with questions related to censorship response but a crucial problem is evaluating the resources and the utility function of censors.
- Ethnological approaches may be useful in studies to integrate into communities affected by flooding (effective information operations) or fear-based censorship. However such studies may not provide quantitative data about successes in circumvention efforts. Safety training and education might provide opportunities for these ethnographic studies
- Policy and the threat of regulatory responses may be the most effective tool to ensure "good behavior" by the operators of centralized systems, but it was observed that GDPR has had some counter-intuitive or unintended consequences that in some cases have increased or merely shifted centralization.

Of course systems, machine learning and stylometric techniques, and applied cryptographic approaches should also continue to be pursued in combination with results obtained from these new approaches. Some of these changes (e.g., block-chain, IPv6) may provide opportunities as well as challenges for the private and resilient delivery of content.

4. Long-Term (> 10 years) Significance

We fully expect that network conditions will continue to evolve, as new technologies, applications, and economic and regulatory regimes emerge. These seem likely to result in continued challenges for network anonymity and censorship circumvention, and measurements to understand and evaluate these schemes. Given the fundamental nature of the rights these technologies enable, we expect the problem and the tensions between some approaches to network security and these properties to remain relevant and significant indefinitely.

Appendix: Meeting Notes and Attendees

Challenge 1: How do we get and share data ethically?
- Needed for reproducibility
- What data is available?
- How can we safely measure people?
- IRBs are wildly different in what they will allow
- Feels like any measurement of users may be off the table for some IRBs
- De-anonymization
- Can differential privacy help? Seems unlikely for some low-level timing data, but the census seems to be getting it work for some things.
- Can generative models help? Probably only when you already know what patterns are important.
- Can we use secure enclaves?
- What good models of data sharing exists
- Case study: ISPs want to share data about interconnects
- Censors don't have IRBs. How do we compete?

Challenge 2: How do we evaluate circumvention approaches?
- Need a model of how censors act (Why does Iran censor Telegram but not VPNs?)
- Need a model of how censors will act given a new transport
- Need empirical data, but see Challenge 1
- We have a lack of visibility, blackbox modeling needed
- Are we winning?
- What's winning?
- Costs, economics factor in

More on Censorship and Circumvention
- Encrypting the internet has consequences for freedom on the internet
- Tor pluggable transports
- Stochastic streaming algorithms for circumvention
- Censors using ML, particularly GANs
- Getting volunteers for supporting proxies for circumvention
- Identifying content likely to be censored using linguistic features
- Detecting country-wide blackouts on the internet
- Network measurements
- How firewalls become censorship
- Should we be measuring TCP filters?
- What happens when governments have apps on your hardware?
- Such as using hashes for looking for banned content (it will start with child pornography)

Challenge 3: What are the trends in concentration?
- DNS at CloudFlare
- Few major network providers
- CDNs
- Tor exits are mostly hosted at a small number of places
- End-to-end correlation attacks getting easier when there's fewer end points
- Lots of Tor traffic uses Google DNS
- A small number of browsers are in use. What if one goes rogue?
- Centralized systems become fragile
- How does concentration relate to filtering?
- Does it matter if a few companies can turn off the internet?
- Companies are driven by economics, but fear of regulation seems to keep them in check
- Jigsaw used VPNs and Digital Ocean. Google's attempt to get rid of Meek?
- On the other hand, the internet seems to be breaking apart into national internets. What does that mean for censorship?

Challenge 4: How do control disinformation without becoming censors?

- Content moderation, Section 230, SESTA
- Flooding and fear, not just friction
- Hard study to flooding/disinformation on WhatsApp
- Can we measure the effects of fake news?
- We can see the effects of pump and dumps on markets
- Who gets to decide what's OK?
- How about provenance? Breaks anonymity.
- Anonymity is context specific
- Social sciences for understanding communities
- Websites blocking users

Privacy

- The privacy threat is no longer ISPs but websites
- Threats exist at all levels of the stack
- Authorship attribution
- Metadata
- Blockchain for anonymous communication
- Privacy while circumventing censorship
- Using public keys for anonymous reputations

Need for SBE

- A lot of this comes down to economics, already noted that censorship often boils down to costs
- Already noted need for social understanding of communities surrounding disinformation
- How does GDPR affect ads? Competition? Helps large firms that can get consent?
- SBE needed, but they don't want to just deploy their methods to solve our problems. How can we make it interesting for them?

Misc.

- Training people in Africa on how to securely use the internet
- Carrier-grade NATs pose problems for tracking the sources of abuse
- Cybersecurity policy
- Cloud security
- IoT
- Cell phone networks and controlling who can call you (providers know who you are but can't share it due to laws)

Attendees

FIRST	LAST	Email
Mark	Allman	mallman@icir.org
Paul	Barford	pb@cs.wisc.edu
Richard	Brooks	rrb@g.clemson.edu
Ken	Calvert	kcalvert@nsf.gov
Paulo André	Da Silva Gonçalves	pasg@cin.ufpe.br
Nick	Feamster	feamster@uchicago.edu
Stephen	Hall	stephen.hall.ctr@darpa.mil
Nick	Hopper	hoppernj@umn.edu
Amir	Houmansadr	amir@cs.umass.edu
Rob	Jansen	rob.g.jansen@nrl.navy.mil
Patrick	Juola	juola@mathcs.duq.edu
Aniket	Kate	aniket@purdue.edu
Susan	Landau	susan.landau@tufts.edu
Chris	Leberknight	leberkc@gmail.com
Anita	Nikolich	anita.nikolich@gmail.com
Indrajit	Ray	iray@nsf.gov

Brad	Reaves	bgreaves@ncsu.edu
Vyas	Sekar	vsekar@andrew.cmu.edu
Micah	Sherr	msherr@cs.georgetown.edu
Craig	Shue	cshue@cs.wpi.edu
Michael	Tschantz	mct@icsi.berkeley.edu
Vinod	Yegneswaran	vinod@csl.sri.com