My research employs empirical methods to study large-scale Internet attacks. As the Internet's user base and criticality of online services continue to expand daily, nation-state adversaries like Internet censors are increasingly monitoring and restricting Internet traffic. Censors perform large-scale *tampering* attacks seeking to prevent users from accessing specific online content, compromising Internet availability and integrity. In recent years, we have witnessed recurring censorship events affecting Internet users globally, including in my home country, India, with far-reaching social, financial, and psychological consequences. However, characterizing tampering attacks is an extremely challenging problem especially at the global scale, given intentionally opaque practices of censors, varying tampering mechanisms across networks, sparse ground truth, and safety risks in collecting data.

My goal is to build frameworks that promote a more globally accessible Internet. I build empirical techniques, platforms, and data analysis methods to study tampering on the global scale [8 (SIGCOMM'23), 9 (CCS'20), 11 (FOCI'23)], introduce approaches to study key network events rapidly as they occur [6 (USENIX SEC'23), 7 (IMC'20)], and investigate the network technology that enables tampering [10 (NDSS'20a), 12 (CoNEXT'22)]. While regional access is governed by local laws and jurisdictions, we see today that unregulated Internet restrictions are leading to the emergence of increasingly isolated online bubbles. I believe that we as educators and pioneers must strive to create technology that ensures users have access to information when they need it. To that end, my career vision is to build frameworks, tools, and policies that enable informed decision-making regarding an accessible, global Internet. This includes monitoring of superfluous restrictions imposed by both local and foreign actors, the creation of technologies that facilitate transparency and user choice, and the establishment of standards and policies regarding delivering varied content to different users.

Overall, I have contributed to 13 peer-reviewed publications in top-tier venues in Computer Security and Networking, including USENIX Security, ACM Conference on Computer and Communications Security (CCS), Network and Distributed System Security Symposium (NDSS), ACM SIGCOMM, and ACM Internet Measurement Conference (IMC). I have been recognized as a Rising Star at the Free and Open Communications on the Internet (FOCI) 2023, and I was awarded the IRTF Applied Networking Research Prize 2023. I have also received funding through the Open Technology Fund Information Controls Fellowship in 2021 and the University of Michigan Rackham Predoctoral Fellowship in 2023. My work has produced one of the biggest active censorship measurement platforms, the Censored Planet Observatory, and has protected freedom on the Internet for millions of users through enabling policy change and preventing large-scale attacks on end-to-end encryption [2].

Global Detection of Connection Tampering

My current research focuses on obtaining a global, data-driven view of connection tampering phenomena such as Internet censorship. Previously established censorship measurement platforms utilize volunteer networks of end-user devices [3]. While these platforms have been successful in specific case studies on censorship, the need to recruit volunteers safely limits the *scale, coverage, and continuity* of measurements, and the resulting data tends to be sparse and ill-suited for studying tampering globally and longitudinally. I address these challenges by building techniques that measure tampering at scale without requiring explicit end-user participation.

Remote Measurement Techniques: A promising research direction is measuring tampering *remotely*, sending thousands of measurements to public Internet infrastructure to deduce tampering through network side-channels, removing the need for volunteer measurements. My work has introduced *Hyperquack*, a remote measurement technique that measures HTTP GET request- and TLS SNI-based tampering through measurements to large organizational web servers [10 (NDSS'20a)]. Hyperquack uses the insight that web servers frequently send the same error response (e.g. `404 Not Found`) when responding to a request for a domain that the web server does not host. Thus, when Hyperquack requests a blocked domain, a response from the server that differs from this error is indicative of tampering. Because of the abundance of web servers, including those hosted by ISPs themselves, Hyperquack is able to measure HTTP(S) tampering across tens of thousands of public organizational vantage points globally, without sending measurements from or to end-user devices.

The Censored Planet Observatory: I have advanced the state of the art in censorship measurement through my efforts as the lead researcher of *Censored Planet observatory*[1], a global, longitudinal censorship measurement platform that actively and remotely monitors Internet censorship in more than 200 countries [9 (CCS'20)]. While standalone remote measurement tools like Hyperquack face coordination challenges that hinder sustained longitudinal monitoring, the modular design of the Censored Planet Observatory strategically selects, schedules, and runs continuous remote measurements to thousands of vantage points on various protocols such as TCP/IP, DNS, and HTTP(S). The observatory complements volunteer-based platforms by offering wide coverage, running measurements to more than 90,000 vantage points distributed across a median of 8 Autonomous Systems (ASes) per country. Since our launch in 2018, the observatory has pioneered the censorship measurement space and has collected and published more than 80 terabytes of 65 billion data points which have enabled numerous studies on censorship [5 (NDSS'20b), 13 (PETS'23), 15 (FC'21)]. I have enabled dozens of researchers, technologists, journalists, and Internet freedom advocates to monitor Internet censorship using data from the Censored Planet observatory.

My work has also introduced new methods in statistical data analysis and machine learning to accurately analyze large-scale empirical data [9 (CCS'20), 11 (FOCI'23), 14 (NDSS'24)]. In a paper published at ACM CCS 2020, I analyzed 20 months of Censored Planet observatory data by modeling it as a time series using a representative metric for country-level censorship [9 (CCS'20)]. The metric takes into account the heterogeneity in censorship policies between different ASes within the country. I then developed and applied a

[1] `https://censoredplanet.org`

Bitmap-based anomaly detection technique that enables the detection of key censorship events through identification of large changes in the censorship metric over time. Such detection of censorship was previously challenging because of the cumbersome manual effort required for data cleaning and analysis. The anomaly detection technique enabled the automated detection of 15 prominent censorship events, such as the blocking of several social platforms following periods of unrest in Sri Lanka, Venezuela, Zimbabwe, and Sudan. Additionally, the study explored censorship trends over time using the Mann-Kendall statistical test, which showed increasing censorship activity in more than a hundred countries, particularly using DNS and HTTPS blocking methods. More recently, I have contributed to follow-up research that models censorship data as decision trees and uses unsupervised cross-clustering of decision trees to detect prominent events [14 (NDSS'24)]. I have also helped to establish best practices in censorship data analysis and build analysis pipelines and dashboards for efficient exploration of censorship data [11 (FOCI'23)]. **In addition to 7 publications in top-tier security venues such as NDSS and ACM CCS, the above body of work has been featured in more than 80 news articles in venues such as the Associated Press, CPJ, MIT Technology Review, and the Financial Times.**

Passive Measurements: While my remote measurement work involves *actively* sending probes to remote destinations, I have also developed techniques to *passively* measure global connection tampering using large-scale network taps. I worked with Cloudflare, a large content delivery network, to develop a completely passive methodology to measure connection tampering from the traffic received at Cloudflare's network edge [8 (SIGCOMM'23)]. This enables visibility into networks where remote vantage points are difficult to obtain for active measurements, such as in mobile networks. Using manual investigation and insights from previous work, I developed a set of connection tampering *signatures*, TCP packet sequences that are indicative of middlebox interference of client network connections. For instance, the Chinese firewall system sends multiple RST-ACK packets to both the client and the server when observing a request for censored content. Applied to global traffic arriving at Cloudflare, the signatures are able to both validate observations about censorship from previous work, but also uncover new instances of tampering in countries such as Peru that have received little attention in active measurement work. More importantly, the passive analysis helps detect tampering as experienced by real-world clients, which reveals the magnitude of real-world tampering across time. The results revealed the heterogeneity of tampering policies among different ASes within countries, and how connection tampering increased in Iran following protests in September 2022. **Cloudflare has deployed this work in a first-of-its-kind passive observatory of connection tampering.**

Next Steps: As a faculty member, I will continue to drive this thrust forward by establishing systems that gain insight into previously underexplored forms of Internet restrictions. For instance, I will explore partial content blocking, where specific segments of a web service are restricted or throttled as opposed to complete blocking of the service. I will overcome challenges such as lack of vantage points and absence of definitive ground truth by expanding existing measurement networks such as the Censored Planet observatory.

Rapid Study of Evolving Network Events with New Threat Models

My current research examines evolving network events during momentous geo-political occurrences, like conflicts, that lead to increased restrictions on Internet freedom. My work systematically documents, reacts to, and defends against emerging threat models and new forms of Internet restrictions through the rapid development and deployment of new empirical methods.

Network Changes in Russia during the Conflict with Ukraine: I co-led a multi-perspective study into new types of network restrictions that emerged in Russia after the escalation of the conflict with Ukraine in February 2022 [6 (USENIX SEC'23)]. Russia's invasion of Ukraine was immediately followed by a torrent of increased censorship, geoblocking, circumvention usage, and the rise of a domestic certificate authority in Russia. Using new measurement methods and data synthesized from nine independent data sources, our study characterized these emerging threats to Internet freedom.

Due to increased sanctions against Russia and Russia's isolation of their own network space, the conflict led to increases in *geoblocking*, a phenomenon where specific user populations are denied access by content providers based on their location. I developed an open-source measurement tool, *GeoInspector*, that uses DNS, TCP, and HTTP measurements and fingerprints to identify geoblocking. GeoInspector uses measurements from within the region of interest, as well as control measurements from other regions, to identify cases where content providers block access to certain users. Through measurements in 16 countries from March–May 2022, GeoInspector detected increased geoblocking by both Russian as well as foreign content providers. Around 25% of Russian government domains tested blocked all visitors outside of Russia, and another 20% of government domains blocked access from all countries except Russia and Kazakhstan. On the other hand, 159 popular domains hosted on foreign content providers geoblocked Russian users specifically, and another 67 popular domains were unavailable in Russia and Kazakhstan. My analysis of GeoInspector data showed that government, education, and even news websites are geoblocking Russian users, isolating them further and creating an information bubble.

In addition to geoblocking, our multi-perspective study also investigated the rise of a new domestic certificate authority, which raised concerns of surveillance. We also investigated the increase in censorship of popular domains such as Twitter and BBC through data from the Censored Planet observatory, OONI, Routeviews, and IODA, and the increase in circumvention tool usage through data from Tor, Psiphon, and VPN providers. **Our findings serve as a cautionary tale for Internet freedom. The landscape of Internet restrictions is rapidly changing, with an array of private actors joining a growing number of government actors in implementing restrictions.**

HTTPS Interception in Kazakhstan: In July 2019, the Republic of Kazakhstan launched a large-scale HTTPS Interception Man-in-the-Middle (MitM) attack on encrypted traffic in the country, after instructing citizens to install and trust a government-issued root "security" certificate that would give them the power to intercept and read encrypted traffic. The interception, described as a pilot attempt, covered large portions of the country's traffic intermittently until being shut down in August 2019. I studied the interception while it was active, through both remote and in-country measurements [7 (IMC'20)]. My investigation involved performing remote TLS handshake measurements to 6,736 TLS hosts in Kazakhstan, with the TLS Server Name Indication (SNI) value set to various domain names to identify which names triggered the interception. The study also used TTL-limited network probes to identify the network location of the interception system, and recorded longitudinal measurements to track the attack over time. The measurements identified the interception in probes to 7%–24% of the TLS hosts, and pinpointed the interception system within AS 9198 (Kazakhtelecom), the biggest and state-owned ISP in Kazakhstan. The interception specifically targeted 37 domains, among which were domains belonging to popular social media and content platforms such as Google, Facebook, Twitter, and mail.ru. The set of targets suggests that the government's actions were motivated by surveillance, rather than increased security as was officially claimed.

The interception attack represents a completely new threat model, where powerful adversaries are able to defeat the security guarantees provided by HTTPS. Kazakhstan's national-level HTTPS interception sets a dangerous precedent—not only for Kazakhstan, but for all governments and other powerful actors that want to exercise more control over users' Internet traffic. **Based on my findings, two major browser vendors, Mozilla Firefox and Google Chrome, completely blocked the use of the government-issued root certificate [2]. My work was highlighted by more than 50 media organizations worldwide, such as BBC, EFF, Forbes, and Ars Technica.**

Next Steps: I will continue to research novel attacks by nation-state adversaries that pose new challenges to Internet security. For instance, I intend to conduct investigations scrutinizing the threat models and underlying assumptions inherent in security protocols and technologies like TLS, encrypted DNS, and VPNs. My objective is to determine whether these safeguards are vulnerable to threats posed by highly powerful adversaries. Additionally, I intend to explore the involvement of private actors such as content providers in imposing user restrictions, and develop incentives to discourage such practices.

Investigating Network Devices Performing Deep Packet Inspection

My current research investigates the proliferation of powerful network devices that can inspect, filter, and tamper with large-scale traffic using Deep Packet Inspection (DPI) technology. General-purpose solutions for studying DPI devices that perform blocking can encourage more accountability and oversight for both device manufacturers and the actors that deploy them. However, collecting features about DPI devices at scale is challenging due to the cumbersome manual effort involved in collecting device fingerprints, the large variety of devices and tampering methods, and the lack of transparency by network device vendors. The inability to monitor the proliferation of DPI devices broadly prevents researchers and regulators from efficiently discovering and responding to the misuse of these dual-use technologies.

Fingerprinting DPI Devices based on Blockpages: A key feature of many DPI devices tampering with application-layer DNS and HTTP traffic is that they frequently return *blockpages*, web pages that indicate the presence and sometimes even the reason and entity behind the tampering. I developed a semi-automatic detection framework, *FilterMap*, that identifies deployments of network devices using the blockpages that they return [10 (NDSS'20a)]. FilterMap takes in millions of censorship measurements collected by observatories like Censored Planet and applies image clustering methods to automatically group blockpages that are returned by vendors or deployers of the same DPI device. Filtermap detected 70 unique clusters of blockpages corresponding to DPI device deployments in more than a hundred countries—some at the national or ISP level and others at the corporate or institutional level—among which were products from well-known commercial manufacturers including Fortinet and Cisco. The use of these technologies for the purpose of blocking across many countries emphasizes the need for regulators to increase visibility into the growth of DPI devices.

Novel Methods for Identifying DPI Device Features: I co-led the development of new empirical methods to identify the location and rules of DPI devices [12 (CoNEXT'22)]. The study introduces a censorship traceroute tool, *CenTrace*, that uses TTL-limited measurements to automatically identify where DPI devices are located in the network, without any information about what devices are used or how the blocking is implemented. CenTrace accounts for a wide variety of DPI device behaviors, such as off-path blocking and IP header modification. Through 12,600 traceroutes in four countries—Azerbaijan, Belarus, Kazakhstan, and Russia—CenTrace found evidence of DPI device deployments both in client as well as upstream ISPs, the latter affecting a larger portion of networks. Interestingly, CenTrace even found cases of device deployments in one country (Russia) affecting traffic destined to another country (Kazakhstan), revealing the importance of studying the network location of censorship in a field where the attribution of blocking is critical but challenging.

I also built a censorship request fuzzing tool, *CenFuzz*, that identifies the rules and triggers of DPI device deployments, with the insight that devices from the same vendor or deployer would exhibit similar behavior to fuzzed requests. CenFuzz consists of 16 HTTP request and eight TLS Client Hello fuzzing strategies that automatically attempt to evade detection by the DPI device. Clustering the result of fuzzing strategies revealed that devices manufactured by the same vendor form tight clusters. The study paves the way for new approaches to fingerprint network devices, and open-sources tools that enable the continued monitoring of global DPI device deployments. **My work enables researchers, journalists, policymakers, and Internet freedom advocates to demand accountability from device manufacturers and the authorities that deploy them. For my work's contribution to the networking community, I was awarded the IRTF Applied Networking Research Prize 2023.**

Next Steps: Moving forward, my research interests will encompass auditing of middlebox blocklists and the application of privacy-preserving technologies like Zero Knowledge middleboxes. Additionally, I will conduct studies on the use of machine learning-based models by these devices for content blocking, and develop adversarial examples that can help with circumvention.

Future Research Agenda

Going forward, I will focus on advancing the empirical study of large-scale tampering attacks and building privacy enhancing technologies that address disparities in Internet access. In addition to expanding my current research lines as described above, I will (1) develop novel systems that can optimize measurements and automatically react to changes in measurement environments; and (2) investigate the impact of Internet restrictions, their underlying causes, and the users that are impacted.

Building *Intelligent* Measurement Platforms: My work has highlighted that emerging threat models in network security necessitate improvements in measurement techniques. For instance, reacting to new forms of tampering or changes in censorship policies currently require both information from on-the-ground collaborators as well as cumbersome manual effort in changing measurement configurations. This often leads to researchers employing a sub-optimal, uninformed exploration approach that does not efficiently leverage the limited measurement resources and time at their disposal. Ultimately, this results in incomplete gathering of data during critical periods.

I propose building *intelligent* measurement platforms that can optimize measurement of connection tampering. Building on my expertise in developing machine learning models for censorship data [10 (NDSS'20a), 14 (NDSS'24)], I will develop Reinforcement Learning (RL) models that perform measurements resulting in increased utility based on knowledge gathered from previous measurements. I will model tampering measurements as an environment where an RL agent can take actions such as adding new websites to test based on features such as the category of the website. I will design reward functions that expedite and automate the discovery of changes in the tampering environment. I envision constructing a multifaceted system featuring multiple RL models tailored to diverse use-cases, such as performing measurements on additional protocols during times of increased tampering. I believe this system will significantly advance the current state of the art, enabling researchers to respond more effectively to tampering events and conserve valuable resources. Moreover, the intelligent platform can automatically alert researchers when new types of content are blocked.

In addition to advancing tampering measurements, I will expand the capabilities of intelligent platforms to efficiently measure the exploitation of security vulnerabilities on a global scale. My experience in working with the Citizen Lab has revealed that detecting network security exploits, especially targeted attacks, is a labor-intensive, manual procedure that faces challenges in scalability due to the absence of reliable ground truth and measurement capabilities. While tools like ZMap offer possibilities in this regard, current approaches are still limited by uninformed exploration and require manual intervention for targeted attacks. I will develop RL models that utilize remote measurements and security exploit fingerprints to automatically collect data on a large scale and adapt to changing environments. Such intelligent measurement platforms can provide the research community with sustainable data for years to come.

***Why* Internet restrictions and *Who* is Affected?** So far, I have scrutinized *how* Internet restrictions are implemented [6 (USENIX SEC'23), 7 (IMC'20), 10 (NDSS'20a)], *where* it occurs [12 (CoNEXT'22)] and *what* content and networks are affected over time [4 (FOCI'21), 8 (SIGCOMM'23), 9 (CCS'20)]. However, the community has yet to investigate *why* certain content is blocked, and *who* is affected by Internet restrictions. In my future research, I will conduct studies on (1) the *real-world impact* of Internet restrictions, which includes interdisciplinary research on the users that are affected by restrictions such as censorship; and (2) the *reasons* why tampering occurs, encompassing investigations into the characteristics of online content that make them more susceptible to tampering.

The lack of knowledge on the range of users affected by Internet restrictions represents a critical gap, as advocacy based on the findings of measurement studies is currently limited to what content is affected, and not *how many* or *what type* of users are experiencing tampering. Understanding the impact of restrictions on users requires a multi-perspective investigation combining active measurements, human-centric analysis, and data collection from content providers. Each of these approaches provides a critical view of the ecosystem, and I will use my expertise in these areas to overcome unique challenges such as the ethical collection of user data, privacy-preserving means of server-side data sharing, and Internet-scale measurement collection. For example, my expertise with DPI devices will aid in finding the number of users that are affected by a certain DPI device through multi-client IP address identification and network population size estimation. Moreover, my experience in collaborating with large content providers has helped me understand the nuances of server-side data sharing, and I will build standards and incentives for data sharing that preserve the privacy of providers, clients, and customers. Ultimately, understanding the real-world impact of Internet restrictions will enable data-driven advocacy for Internet freedom.

Analyzing the reasons behind Internet restrictions requires studies into the features of online content and traffic that may be targeted. For instance, do online communities that provide users with anonymity face more tampering? Exploring such a research question requires a systematic study that collects data about Internet content that is affected and identifies patterns, both through qualitative and quantitative methods. I will build frameworks that can scale up such investigations through content analysis and web measurement. I believe these frameworks can help providers of blocked services understand the underlying causes and intention behind tampering.

In the longer term, I will incorporate surveys and qualitative interview data from users affected by Internet restrictions to gain more insight into the *human and economic cost* of such restrictions. For example, although there is anecdotal knowledge that restrictions like Internet shutdowns lead to significant economic losses, there is little research into the practical impact on day-to-day business operations, which I will explore through these studies. However, recruiting users for participation involves significant safety risks, and I will draw on my experience in working with such users in the past to design studies that minimize risks through engaging with the community and developing safety standards. Moreover, I will also conduct a range of studies into the security, privacy, and usability features provided by alternate local platforms and software, such as state-approved online platforms and browsers. I have observed previously that these state-approved platforms can gain traction due to Internet restrictions [1 (USENIX SEC'22), 6 (USENIX SEC'23)]. Understanding the security and privacy properties provided by these platforms is imperative to safeguard user data from monitoring and tampering. Through these lines of research, I ultimately want to build and foster an open and safe Internet for all users.

References

[1] R. Kumar, A. Virkud, **R. Sundara Raman**, A. Prakash, and R. Ensafi. A Large-scale Investigation into Geodifferences in Mobile Apps. In *USENIX Security Symposium*, 2022.

[2] Mozilla. Mozilla takes action to protect users in Kazakhstan. The Mozilla Blog, August 21, 2019. `https://blog.mozilla.org/blog/2019/08/21/mozilla-takes-action-to-protect-users-in-kazakhstan/`.

[3] OONI: Open Observatory of Network Interference, 2012. `https://ooni.org/`.

[4] R. Padmanabhan, A. Filastò, M. Xynou, , **R. Sundara Raman**, K. Middleton, M. Zhang, D. Madory, M. Roberts, and A. Dainotti. A Multi-Perspective View of Internet Censorship in Myanmar. In *ACM SIGCOMM 2021 Workshop on Free and Open Communications on the Internet (FOCI)*, 2021.

[5] R. Ramesh, **R. Sundara Raman**, M. Bernhard, V. Ongkowijaya, L. Evdokimov, A. Edmundson, S. Sprecher, M. Ikram, and R. Ensafi. Decentralized Control: A Case Study of Russia. In *Network and Distributed System Security Symposium (NDSS)*, 2020.

[6] R. Ramesh*, **R. Sundara Raman***, A. Virkud, A. Dirksen, A. Huremagic, D. Fifield, D. Rodenburg, R. Hynes, D. Madory, and R. Ensafi. Network Responses to Russia's Invasion of Ukraine in 2022: A Cautionary Tale for Internet Freedom. In *USENIX Security Symposium*, 2023.
** indicates joint authorship.*

[7] **R. Sundara Raman**, L. Evdokimov, E. Wustrow, A. Halderman, and R. Ensafi. Investigating Large Scale HTTPS Interception in Kazakhstan. In *In ACM Internet Measurement Conference (IMC)*, 2020.

[8] **R. Sundara Raman**, L.-H. Merino, K. Bock, M. Fayed, D. Levin, N. Sullivan, and L. Valenta. Global, Passive Detection of Connection Tampering. In *ACM SIGCOMM*, 2023.

[9] **R. Sundara Raman**, P. Shenoy, K. Kohls, and R. Ensafi. Censored Planet: An Internet-Wide, Longitudinal Censorship Observatory. In *In ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2020.

[10] **R. Sundara Raman**, A. Stoll, J. Dalek, R. Ramesh, W. Scott, and R. Ensafi. Measuring the Deployment of Network Censorship Filters at Global Scale. In *Network and Distributed System Security Symposium (NDSS)*, 2020.

[11] **R. Sundara Raman**, A. Virkud, S. Laplante, V. Fortuna, and R. Ensafi. Advancing the Art of Censorship Data Analysis. In *Free and Open Communications on the Internet (FOCI)*, 2023.

[12] **R. Sundara Raman***, M. Wang*, J. Dalek, J. Mayer, and R. Ensafi. Network Measurement Methods for Locating and Examining Censorship Devices. In *In ACM International Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, 2022.
** indicates joint authorship.*

[13] E. Tsai, D. Kumar, **R. Sundara Raman**, G. Li, Y. Eiger, and R. Ensafi. CERTainty: Detecting DNS Manipulation using TLS Certificates. In *Privacy Enhancing Technologies Symposium (PETS)*, 2023.

[14] E. Tsai, **R. Sundara Raman**, A. Prakash, and R. Ensafi. Modeling and Detecting Internet Censorship Events. In *Network and Distributed Systems Security Symposium (NDSS)*, 2024.

[15] A. Vyas, **R. Sundara Raman**, N. Ceccio, P. M. Lutscher, and R. Ensafi. Lost in Transmission: Investigating Filtering of COVID-19 Websites. In *Financial Cryptography and Data Security (FC)*, 2021.