AMIR HOUMANSADR RESEARCH STATEMENT

The Internet has become ubiquitous across the globe with more people getting connected to it, at more places, and through increasingly varied network-enabled devices. Unfortunately, Internet users face constant threats to their online security and privacy: repressive regimes deprive them of open access to the Internet, corporate networks monitor their online behavior, advertising companies collect and share their private information, and cybercriminals hurt them financially through security breaches. I am interested in *building secure and privacy-preserving tools for Internet communications*. Such tools can be deployed not only by dissidents and whistleblowers, but also by ordinary Internet users on a daily basis.

More specifically, I assess the security and privacy provided by existing network protocols and services, and propose design adjustments or clean-slate architectures to retain users' security and privacy. To this end, I combine the development of practical systems with rigorous theoretical analysis, and incorporate techniques from various disciplines such as computer networking, cryptography, and statistical analysis. The specific problems I have explored in the past include Internet censorship resistance, network traffic analysis, network situational awareness, social network malware, mobile security, and multimedia information hiding. My research has found flaws in popular privacy-preserving tools, and has led to the advent of novel designs to overcome these problems. My work has been publicized in the media through interviews, and I received the *Best Practical Paper award* from the IEEE Symposium on Security & Privacy 2013.

1 The Cyberspace Battle for Information: Combating Internet Censorship

The Internet plays a crucial role in today's social and political movements; democracy and human rights throughout the world critically depend on preserving and bolstering the Internet's openness. To prevent the Internet from achieving its democratizing potential, authoritarian regimes monitor and limit their citizens' access to the Internet. Many systems for circumventing such barriers have been deployed to help the affected users, e.g., HTTP proxies and Tor. For a decade, the Internet community has been part of a cat-and-mouse game between circumvention systems and the censors deploying increasingly advanced approaches. Recently, the balance has, unfortunately, tipped the way of the censors. This is, in part, due to larger investments in censorship by repressive regimes triggered by the Arab Spring movement, globally reduced prices of censorship equipment manufactured by competing networking companies, and recruiting Western-educated technicians and scientists by the censors.

I believe that research on censorship circumvention is one of the most important, influential research topics as it has the potential to *significantly affect billions of Internet users globally*. Also, censorship resistance is a technically demanding, exciting research problem: First, censorship resistance is an inherently interdisciplinary problem, spanning the areas of security, cryptography, and networking. Second, it has intrinsic correspondence with various non-technical sciences such as social sciences, public policy, and political sciences. The failure story of circumvention systems in the past several years testifies the difficulty and complexity of designing and building practical circumvention systems.

I contend that a major pitfall of today's circumvention systems is their inability to remain *unobservable*, e.g., they are easily detectable at the network level. Existing designs mostly rely on ad-hoc heuristics that, when discovered by the censors, are easily exploited to interfere with their operation. I argue that *censorship resistance should be integrated into essential, popular network protocols and services*, like the TLS protocol and Skype, in order to provide *survivable* unobservability. A major challenge to this approach is providing reasonable quality of service to circumvention users, e.g., latencies bearable for interactive web browsing.

Decoy Routing: I invented a novel, clean-slate architecture for censorship resistance known as decoy routing [1]. Traditionally, censorship resistance is deployed at the edges of the Internet, e.g., on computers serving as proxies or volunteer websites relaying traffic. In decoy routing, censorship evasion is shifted from the edges of the Internet to the middle, i.e., to the routers and ISPs. More specifically, our architecture, called Cirripede [1], installs circumvention software on volunteer Internet routers, called decoy routers. To utilize Cirripede, a censored user establishes a TLS connection with an arbitrary, non-blocked Internet destination, which we call the overt destination. The client selects the overt destination such that the Internet route to the destination passes through one of the decoy routers. Through this connection with the overt destination, the client steganographically signals the decoy router on the path to serve the connection as a circumvention connection. The decoy router intercepting this connection proxies it to a covert destination of the client's choosing, e.g., a blocked website.

Decoy routing provides strong censorship resistance properties. First, it is highly unobservable: from the perspective of a censor monitoring the traffic of a decoy routing client, the client appears to be communicating with a non-blocked, legitimate destination, while the client is in fact communicating with forbidden destinations. Second, decoy routing is highly unblockable. In order to block, censors need to modify routing decisions of all of their Internet users so that no traffic is intercepted by decoy routers. However, we demonstrate [2] that this incurs significant monetary costs and collateral damage to the censors, e.g., it significantly increases the latency for all Internet users of the censors, including the lawful, benign users.

How to Survive Protocol Identification: Today's censors are capable of detecting circumvention traffic based on protocol identifiers such as timing patterns and content signatures. This has initiated a whole line of research towards designing circumvention tools that hide protocol identifiers, thus allowing themselves to remain unobservable to censors. *Parrot systems* are a particular class of such designs that have become increasingly popular in recent years. A parrot circumvention system evades protocol identification by imitating popular Internet protocols like Skype and HTTP, e.g., by mimicking their packet timings.

We conducted [3] the first in-depth analysis of parrot circumvention systems by developing a taxonomy of censorship adversaries and by enlisting the technical requirements that a parrot system must satisfy to provide unobservability. We analyzed several parrot systems, recently proposed in top security conferences and deployed by Tor, and demonstrated that *all of these systems completely fail to achieve unobservability*. An important, striking conclusion of our study is that "unobservability through imitation" is *fundamentally flawed*. To imitate a network protocol like Skype, a parrot circumvention system needs to satisfy a daunting list [3] of imitation requirements; censors, on the other hand, only need to find a few discrepancies, thus tipping the balance dramatically in the censors' favor.

Motivated by the lessons of our analysis [3] of parrot systems, we suggest an alternative approach for unobservable circumvention, *hide-within*. Similar to the parrot systems, a hide-within system aims to make its traffic look like popular Internet protocols such as Skype. By contrast, instead of imitating, a hide-within system *runs* the target protocol and embeds circumvention traffic into the higher layers of the network protocol stack. We have designed and prototyped three hide-within circumvention systems: FreeWave [4], SWEET [5], and CloudTransport [6]. For instance, FreeWave [4] tunnels circumvention traffic into VoIP calls by encoding network packets into acoustic signals and feeding them to third-party VoIP client software. Our hide-within designs [4–6] provide strong unobservability by tunneling traffic into the genuine runs of popular Internet protocols such as Skype and SMTP. As hide-within traffic is *entangled* with the traffic of benign Internet users, blocking such services will disrupt benign users as well, causing the censors significant collateral damage. Hide-within designs are also *survivable*: unlike other circumvention systems like Tor, identifying a single circumvention client does not make it easier for the censors to identify other clients.

Impact: I proposed Cirripede [1], one of the three concurrent works that invented decoy routing circumvention. Decoy routing has received significant attention from the media and the research community; several academic and industrial research groups, e.g., Raytheon BBN, are actively working towards building a commercial system based on this concept. Our analysis of parrot circumvention systems [3] won the **Best Practical Paper award** at the IEEE Symposium on Security & Privacy (Oakland) 2013, one of the most prestigious venues in the area of security and privacy. My work on censorship resistance has been taught in graduate security classes at several universities, including UC Berkeley, Princeton, and the University of Waterloo. Our hide-within approach has been welcomed by the community and some of our designs are under evaluation and active development in real-world systems. I was interviewed by the College of Natural Sciences newsletter at UT Austin about my research on censorship resistance. Finally, I was invited to write a book chapter [7] on network protocol obfuscation.

2 Chasing Packets Across the Internet: Novel Techniques for Traffic Analysis

Cybercriminals use encryption and other evasion techniques to conceal the malevolent nature of their attack traffic. This has made it extremely hard for network defenders to respond to cyber threats in a timely manner. In particular, cybercriminals commonly relay their attack traffic through *stepping stones*, e.g., previously compromised machines or public proxies, to disguise their identities by hiding the origin of their traffic. This highlights the need for efficient mechanisms that *link* network packets even after getting encrypted

and relayed. Such mechanisms are known as *traffic analysis* and work by statistically correlating traffic characteristics such as packet timings. In addition to tracking down cybercriminals, traffic analysis has other important security-related applications, e.g., in compromising anonymity networks like Tor.

Traditionally, traffic analysis is performed passively, i.e., flows are linked solely by passive monitoring of network traffic. The main challenge to this approach is not being scalable to large-size networks: for a network with m ingress/egress flows the computation and communication overheads are $O(m^2)$ and $O(m)$, respectively. To address this limitation, an active approach has been proposed for traffic analysis, called *flow watermarking*. Flow watermarking perturbs traffic characteristics of network flows to create a distinct pattern, the watermark, which is then looked for to link flows. This reduces the computation and communication costs to $O(m)$ and $O(1)$, respectively. It also improves linking precision and accelerates detection.

Compromising Marks with Regularity. For a flow watermark to be effective, it needs to be *invisible* [8], i.e., its presence on network flows should be undetectable by third-parties. Early watermarking designs evaluated invisibility only against an adversary who monitors individual network flows. We introduced [9] the *multi-flow attack (MFA)*, a new class of attacks that aggregates the observations from different network flows. We demonstrated that the MFA attack can easily be mounted in nearly all applications of flow watermarking and against *all* previously proposed watermarking techniques. Our MFA attack is able to not only detect watermark presence, but also recover secret watermarking parameters and remove watermarks from their carrying flows. Subsequently, we suggested several countermeasures to the MFA attack [10].

Introducing Non-Blind Architectures. Early flow watermarks suffered from various attacks that targeted their invisibility, and endured a poor performance in linking flows. Such shortcomings mostly stem from their *blind* architectures. In a blind design, watermarking parties exchange no information about the network flows they intercept. Early blind watermarks used large watermark signals to compensate for this, which made them prone to various detection attacks. We were *the first* to propose non-blind flow watermarking to overcome these shortcomings. Our construction, RAINBOW [8], imposes packet delays two orders of magnitude smaller than that of previous designs, and improves linking performance by *several orders of magnitude*. We also used [11] "detection theory" to theoretically analyze non-blind watermarks.

Building Scalable and Efficient Schemes. Despite their better invisibility and performance, non-blind watermarks do not scale well with the size of the network due to their use of cross-communications. We proposed SWIRL [12], a blind, hence scalable, watermark that also provides reasonable invisibility. We argue that SWIRL is the first design to be practical for use in large-scale traffic analysis due to its use of a blind design that is also resilient to known attacks. We also demonstrated that SWIRL can be deployed by Tor relays to protect against a recently discovered attack which can exhaust their bandwidths.

Impact: Since our introduction of the MFA attack [9], resistance to this attack has been considered as a design principle in all new flow watermarking designs. My research has given birth to several new concepts in the field of traffic analysis, most notably non-blind flow watermarking [8, 11] and network flow fingerprinting [13]. My RAINBOW [8] and SWIRL [12] designs improved upon the state-of-the-art in traffic analysis by several orders of magnitude. My work on flow watermarking is well-known in the community, has received more than 170 citations, and has been taught in graduate classes in several universities. I was invited to write a book chapter [14] on flow watermarking, to be printed by Wiley-IEEE Press.

3 Research Agenda

I am interested in building *secure and privacy-preserving communication tools for the Internet*. Fortunately, cybersecurity is extremely relevant to several important areas of computer science and engineering such as big data, cloud computing, mobile computing, networking, and online social networking. *I plan to embrace this exciting opportunity for interdisciplinary research and collaboration.* In the past, I have collaborated with more than 30 co-authors in various disciplines, and I plan to continue this trend.

3.1 Re-Thinking the Internet to Evade Censorship: The current Internet struggles mightily against censorship, which is due to one main reason: it was not designed with censorship avoidance as a primary goal. The networking community is actively researching clean-slate architectures for the next-generation Internet to address existing challenges in security, reliability, and scalability. The leading architectures include NDN, MobilityFirst, NEBULA, XIA, and ChoiceNet, which are based on various principles such

as content-oriented networking, service-oriented networking, etc. I believe that *censorship resistance must cease to be an ad-hoc, add-on component and instead become a first-class property of modern networking infrastructure*. Unfortunately, this is not evident in any of the leading designs named above.

I plan to **investigate whether—and how—specific proposals for the next-generation Internet are vulnerable to censorship**. My initial evaluation suggests that these architectures are in fact *more prone to censorship* compared to today's Internet. In particular, content-centric addressing and content caching, which are deployed by most of these architectures, make censorship technically easier. As these proposals are still in the planning and design phases, the outcome of my research can be used to adjust the technical specifications of these architectures to incorporate censorship evasion. At the same time, I do not advocate basing the entire Internet architecture on censorship resistance. Evasion-centric designs are unlikely to be welcomed by the big Internet operators. ISPs, in particular, are likely to be more interested in adopting network techniques that increase their profits, such as content caching in content-centric networks which could reduce their transit costs. I plan to work on finding the right balance between censorship resistance and other principles, e.g., performance, to motivate the global adoption of such new architectures. I believe that finding the answers to these research problems requires *close collaboration with core network researchers, privacy scholars, and legal experts*.

I also plan to **design censorship-resistant mechanisms for future Internet architectures**. I envisage that censorship will remain a serious challenge in any of the proposed successors to today's Internet. Existing circumvention tools will no longer work in the future Internet, and, more importantly, all proposed architectures are likely to facilitate censorship. One research direction that I plan to pursue in this area is *censorship-resistant content naming*. I plan to investigate several approaches that enable users to request censored content, while preventing censors from filtering out these requests and identifying the users who issued them. Another research direction is developing *content delivery abstractions that are secure against side-channel analysis*. Future Internet architectures will likely be vulnerable to side-channel traffic analysis, and I plan to investigate whether today's techniques work for identifying traffic in privacy-preserving content-oriented networks such as ANDaNA.

3.2 User-Driven Privacy for Internet Services: Popular Internet services, such as Gmail and Facebook, invade users' privacy in many ways. For years, the research community has investigated and built privacy-preserving Internet services, such as distributed social networks, secure email services, and peer-to-peer VoIP software to replace privacy-invading services. Unfortunately, because many of these alternatives disable appealing features provided by their privacy-invading counterparts, privacy-preserving services have largely failed to gain traction in any meaningful way. For instance, numerous privacy-preserving social networking architectures were proposed as a replacement to Facebook, yet Facebook remains as popular as ever.

Instead of designing clean-slate privacy-preserving Internet services, I plan to **build user-driven privacy tools that enable ordinary Internet users to control their privacy while using existing, third-party Internet services**. These tools will provide three important properties. First, their operation should not require any collaboration or support from service providers. For-profit service providers should not be expected to properly deploy additional privacy-preserving measures as such measures often conflict with their business models. Even if they do, users' privacy remains at risk from both security compromises and legal threats. For instance, the FBI forced Lavabit, an email provider, to give away its private SSL keys to access Edward Snowden's communications, thereby also revealing the communications of all 400,000 of its customers. Second, user-driven privacy tools should have minimal dependence on the technical specifications of their target services. By minimizing their dependence, user-driven privacy tools avoid frequent, resource-intensive modifications to comply with software updates of the target services. This would also make the solution more generally applicable, e.g., a privacy tool targeting email communications can be usable for email services provided by various vendors. Third, such tools should be user-friendly and as transparent as possible to the users. This is indeed a critical requirement: multiple studies have shown that average Internet user tends to disregard security for better usability. Towards this goal, I aim to collaborate with researchers of *usable security, human-computer interaction, and social engineering*.

I plan to investigate user-driven privacy tools for *a broad range of Internet services*, including email services, social networking apps, VoIP software, and cloud services. In particular, I have started to design tools for VoIP communications. My proposal provides a universal socket that allows users of various VoIP software (e.g., Skype and Google Talk), with minimal technical expertise, to encrypt their VoIP communications end-to-end without disabling the features of the underlying VoIP software, or degrading call quality.

References

[1] **Amir Houmansadr**, Giang Nguyen, Matthew Caesar, and Nikita Borisov. Cirripede: Circumvention Infrastructure using Router Redirection with Plausible Deniability. In *the* 18^{th} *ACM Conference on Computer and Communications Security (**CCS**)*, 2011.

[2] **Amir Houmansadr**, Edmund Wong, and Vitaly Shmatikov. No Direction Home: The True Cost of Routing Around Decoys. In *the* 21^{st} *Annual Network & Distributed System Security Symposium **(NDSS)***, 2014.

[3] **Amir Houmansadr**, Chad Brubaker, and Vitaly Shmatikov. The Parrot is Dead: Observing Unobservable Network Communications. In *the* 34^{th} *IEEE Symposium on Security & Privacy **(Oakland)***, 2013.

[4] **Amir Houmansadr**, Thomas Riedl, Nikita Borisov, and Andrew Singer. I Want my Voice to be Heard: IP over Voice-over-IP for Unobservable Censorship Circumvention. In *the* 20^{th} *Annual Network & Distributed System Security Symposium **(NDSS)***, 2013.

[5] **Amir Houmansadr**, Wenxuan Zhou, Matthew Caesar, and Nikita Borisov. SWEET: Serving the Web by Exploiting Email Tunnels. *CoRR*, abs/1211.3191, 2012.

[6] Chad Brubaker, **Amir Houmansadr**, and Vitaly Shmatikov. CloudTransport: Practical Unobservable Networking Using Cloud Storage. *Under Submission.*

[7] **Amir Houmansadr**. Obfuscating Network Protocols. In W. Mazurczyk, S. Wendzel, S. Zander, A. Houmansadr, and K. Szczypiorski, editors, *Information Hiding in Communication Networks: Fundamentals, Mechanisms, and Applications*. Wiley-IEEE Press, 2015. To appear.

[8] **Amir Houmansadr**, Negar Kiyavash, and Nikita Borisov. RAINBOW: A Robust And Invisible Non-Blind Watermark for Network Flows. In *the* 16^{th} *Annual Network & Distributed System Security Symposium **(NDSS)***, 2009.

[9] Negar Kiyavash, **Amir Houmansadr**, and Nikita Borisov. Multi-Flow Attacks Against Network Flow Watermarking Schemes. In *the* 17^{th} *USENIX Security Symposium **(USENIX Security)***, 2008.

[10] **Amir Houmansadr**, Negar Kiyavash, and Nikita Borisov. Multi-Flow Attack Resistant Watermarks for Network Flows. In *the* 34^{th} *IEEE International Conference on Acoustics, Speech and Signal Processing **(ICASSP)***, 2009.

[11] **Amir Houmansadr**, Negar Kiyavash, and Nikita Borisov. Non-Blind Watermarking of Network Flows. *IEEE/ACM Transactions on Networking*, 2014.

[12] **Amir Houmansadr** and Nikita Borisov. SWIRL: A Scalable Watermark to Detect Correlated Network Flows. In *the* 18^{th} *Annual Network & Distributed System Security Symposium **(NDSS)***, 2011.

[13] **Amir Houmansadr** and Nikita Borisov. The Need for Flow Fingerprints to Link Correlated Network Flows. In *the* 13^{th} *Privacy Enhancing Technologies Symposium **(PETS)***, 2013.

[14] **Amir Houmansadr**. Digital Watermarking of Network Flows. In W. Mazurczyk, S. Wendzel, S. Zander, A. Houmansadr, and K. Szczypiorski, editors, *Information Hiding in Communication Networks: Fundamentals, Mechanisms, and Applications*. Wiley-IEEE Press, 2015. To appear.