INDRAPRASTHA INSTITUTE of INFORMATION TECHNOLOGY DELHI

Building Performance, Privacy-Enhancing & Blocking-Resistant Communication Systems

Ph.D. Thesis Defense

Piyush Sharma

ON

17th August @ 8:00 PM IST

Advisor: Dr. Sambuddho Chakravarty
Dr. Mukulika Maity

Examiners:

- Dr. Michalis Polychronakis (Stony Brook University, NY, USA)
- Dr. Kent Seamons (Brigham Young University, Utah, USA)
- Dr. Amir Houmansadr (University of Massachusetts, Amherst, USA)

Please visit https://cse.iiitd.ac.in/events-seminars/ for more details

Abstract:

Free and open communication over the Internet is essential for the overall advancement of modern societies. However, there are numerous ways with which malevolent adversaries try to control the flow of information, by either selectively restricting access to content or in extreme scenarios completely shutting down the Internet. Hence, we aim to augment existing research as well as build novel systems to facilitate open communication over the Internet. Thus, in this thesis, we attempt to answer the question — Can we facilitate open access to Internet by building privacy-enhancing technologies that can also provide blocking resistance, good QoS and deployability? To that end, we propose and build solutions for three broad categories of applications. (i) applications for which there does not exist any functional and deployed solutions (e.g., anonymous calling over Internet) (ii) applications (e.g., accessing censored websites) for which there are functional solutions, but those have great scope for improvement (e.g., blocking resistance, performance) (iii) accessing lightweight Internet applications during events of Internet shutdowns (e.g., posting tweet and retrieving news snippets etc.). Firstly, we explore the feasibility of anonymous voice communication over the Internet as it is of great interest to whistleblowers, privacy practitioners etc. Current literature sidesteps existing anonymity systems such as Tor for voice calling, citing the performance issues that would render real-time voice calling unusable over Tor. Thus, novel systems tailored for voice have been proposed (e.g., Herd). However, the newly proposed systems are not functional, and thus there does not currently exist any usable anonymous voice calling system. Hence, we revisit Tor and perform a comprehensive analysis of performing voice calling over it. With the help of nearly half a million voice calls performed over the real Tor network over a year under various setups, we establish that it is indeed possible to perform anonymous voice calls using Tor. Secondly, we look at accessing censored content (e.g., websites, files etc.) over the Internet. There already exist solutions such as VPNs, Tor etc., that help accessing restricted content by routing it through single (or multiple) proxy nodes. The censor blocks such proxies as soon as they are discovered, while the researchers attempt to build systems to evade such blocking, effectively leading to an arms race. Decoy Routing (DR) is an approach to potentially break this arms race, where a network router (known as the decoy router) doubles up as a proxy, instead of the end hosts. This makes it extremely difficult for the adversary to censor, as it can no longer block based on end hosts IP, but require blocking routers which carry a significant amount of innocuous traffic. However, existing DR solutions use commodity servers to act as decoy routers, as traditional routers cannot be programmed to perform the complex operations required for proxying connections.

This leads to (1) poor performance for end users due to handling of ISP scale flows by commodity servers, and (2) privacy violations of oblivious users (non-DR clients) due i to the analysis of ISP flows by third party DR maintainers. We thus propose SiegeBreaker, which overcomes the aforementioned problems (while also providing other salient features) by building a DR protocol with the help of software defined network (SDN) devices. We show that using SDN provides the essential modularity required in DR designs, eventually leading to good performance along with flexibility in trust relationships for providing better privacy guarantees. We implemented and deployed SiegeBreaker using hardware SDN switches and show with the help of extensive evaluation that SiegeBreaker provides performance comparable to direct TCP downloads. However, despite being a promising solution, DR poses deployment challenges as it relies on the ISPs support to help place the DR infrastructure. Thus, on one hand we have solutions that are readily available (VPNs, Proxies etc.) but are easy to block, and on the other hand we have solutions that are hard to block but face deployment challenges (DR). Thus, next, we attempt to build a solution which is readily available to end users and at the same time provides effective blocking resistance without compromising the performance. To that end we propose a new system Camoufler that utilizes the Instant Messaging (IM) apps to transfer censored content. Using IM makes the system immediately usable as these apps are an integral part of the Internet ecosystem. Moreover, using IM apps as-is to transfer censored content makes it difficult for the adversary to distinguish it from normal IM traffic. We implement Camoufler on five popular IM apps (Signal, Telegram, Slack, Skype and Whatsapp) and demonstrate that a user can access censored websites in a few seconds (3.6s on an average for Alexa top-1K websites). Lastly, we attempt to address a recent and an extreme form of censorship i.e., Internet shutdowns. Such an extreme step of complete Internet disconnectivity leads to various problems for users in such regions (e.g., reporting of power failures, immediate access to medical emergency etc.). To the best of our knowledge there do not exist systems that can provide Internet access during shutdowns. Thus, we built a novel system Dolphin that can provide access to basic Internet services such as Twitter, email, accessing news etc. Dolphin utilizes the cellular voice channel to encode the data bits into audio and transfer it via the cellular call. However, the cellular voice channel is unreliable, inherently lossy, insecure and highly bandwidth constrained. We overcome all the aforementioned challenges and build a complete end to end system. We experimentally demonstrate that Dolphin can be used to access emails, tweets and news all within a few minutes (an email of 500 characters was delivered securely and reliably within 3 minutes).