

TARN: A SDN-based Traffic Analysis Resistant Network Architecture

Lu Yu, Qing Wang, Geddings Barrineau, Jon Oakley,
Richard R. Brooks, and Kuang-Ching Wang
Department of Electrical and Computer Engineering
Clemson University, Clemson, SC, USA
{lyu,qw,cbarrin,joakley,rrb,kwang}@g.clemson.edu

Abstract

Destination IP prefix-based routing protocols are core to Internet routing today. Internet autonomous systems (AS) possess fixed IP prefixes, while packets carry the intended destination AS's prefix in their headers, in clear text. As a result, network communications can be easily identified using IP addresses and become targets of a wide variety of attacks, such as DNS/IP filtering, distributed Denial-of-Service (DDoS) attacks, man-in-the-middle (MITM) attacks, etc. In this work, we explore an alternative network architecture that fundamentally removes such vulnerabilities by disassociating the relationship between IP prefixes and destination networks, and by allowing any end-to-end communication session to have dynamic, short-lived, and pseudo-random IP addresses drawn from a range of IP prefixes rather than one. The concept is seemingly impossible to realize in today's Internet. We demonstrate how this is doable today with three different strategies using software defined networking (SDN), and how this can be done at scale to transform the Internet addressing and routing paradigms with the novel concept of a distributed software defined Internet exchange (SDX). The solution works with both IPv4 and IPv6, whereas the latter provides higher degrees of IP addressing freedom. Prototypes based on Open vSwitches (OVS) have been implemented for experimentation across the PEERING BGP testbed. The SDX solution not only provides a technically sustainable pathway towards large-scale traffic analysis resistant network (TARN) support, it also unveils a new business model for customer-driven, customizable and trustable end-to-end network services.

1 Introduction

Internet was not designed with security and privacy in mind and is littered with security holes. The notorious network censorship, BGP hijacking [10, 32], man-in-the-

middle (MITM) attacks, and distributed Denial-of-Service (DDoS) attacks [25, 35], to name a few, pose real threat to network security and user privacy. These attacks all rely on a seemingly innocuous and crucial element of today's Internet Protocol (IP) to identify their victims – the IP addresses of the communicating source and destination. At domain-level routing, network operators rely on IP prefix-based BGP protocol for inter-domain routing. And what's worse, IP addresses carried in packet headers are always in clear text.

Consequently, national firewalls easily restrict inhabitants' access to information by simply blocking the IP addresses of those websites [8]. While encryption and virtual private networks (VPN) provide some protection, some countries use traffic analysis to further track citizens that access opposition web sites. Similarly, victims of DDoS attacks and MITM attacks have nowhere to hide due to the static binding between the hosts and their uniquely identified IP addresses.

The proposed work explores an alternative end-to-end network architecture called the Traffic Analysis Resistant Network (TARN) that fundamentally removes the vulnerability due to fixed IP addresses of communication sessions by using, instead, multiple short-lived, seemingly random IP addresses. By breaking the static binding between hosts and IP addresses, TARN makes it impossible for attackers to detect communication sessions by observing the source and destination addresses. While in theory traffic analysis methods can analyze other properties of network traffic, e.g., protocol types [9], packet sizes and inter-packet timing [33], to infer the nature of communication, IP addresses are undeniably the prime references used by attackers today. We recognize TARN induces disruptive changes to multiple aspects of the Internet, ranging from domain name lookup to packet forwarding; an incremental approach to gradually introduce TARN into Internet is hence critical. In this paper, we give an overview of three implementation strategies that realizes TARN at different scopes and discuss their implications.

TARN would not have been practical without two recent technology trends – software defined networking (SDN) and cloud-based virtual network functions (VNF). With SDN, end-to-end communication sessions can be identified and modified (in TARN’s case, the IP addresses) with fine grain packet header fields at flexible network locations. With cloud-based VNF, TARN functions can be scalably instantiated driven by user and application demands. As discussed in section 3, besides the option of running TARN at an end host, a cloud-based implementation of TARN, which we call a software defined network exchange (SDX), can be operated by a provider much like the Internet exchange providers today with advanced, highly customizable services for communication sessions. We envision SDXs to offer end users a logically centralized portal to request and configure services with user-chosen attributes, with the actual TARN packet handling to be executed across a distributed network of SDX sites.

The remainder of this paper is organized as follows. Section 2 provides related work in the area of anonymity network and a brief introduction of SDX. In section 3, we introduce the TARN architecture and the implementation strategies. Early experimentation results are discussed in section 4. The paper concludes in section 5 with potential challenges and the future work.

2 Related Work

Various proxy services such as The Onion Routing (Tor) [7], Psiphon [5], and I2P [1], make traffic analysis difficult by inserting intermediate nodes between packet source and destination. The extent to which they can protect communication rely on such intermediate nodes’ ability to remain unknown to adversaries. With their IP addresses being fixed, their success is limited to the time it takes for adversaries to detect them and start analyzing or blocking communications going through them.

Researchers have found the increased latency due to the use of intermediate nodes in such solutions to be a major impediment to user acceptance [26]. I2P is not very widely used. Tor is commonly attacked by countries finding Tor entry nodes and blocking their IP addresses. This can include active probing of nodes suspected of providing entry to the Tor network [29]. Tor is mostly blocked by Chinese government. One Tor initiative avoids traffic blocking by using “pluggable transports” to obfuscate the traffic pattern for the first hop of the proxy. The issue with the traffic obfuscation techniques is that they only obfuscate communication patterns, but not the end hosts [19]. Another concern with proxy networks is that users must trust the intermediate nodes, any of which could be compromised to perform a MITM attack.

Decoy routing, a new alternative, inserts steganography

information into the header of SSL packets which is recognized by “decoy routers” in the backbone, who re-route packets to their intended destination [22]. One drawback is that nodes currently have to attempt connections at random until they find the decoy routing node.

There are other recent SDN-based solutions, e.g., the OpenFlow-based random host mutation (OFRHM) [21], aiming to use OpenFlow to randomly mutate source host IP addresses to avoid unwanted traffic analysis. However, OFRHM uses random addresses within the same subnet of a source node.

The notion of SDX was first described in 2013 by Feamster, et al. [18] as a Software Defined Internet Exchange, applying SDN at Internet exchange points (IXP) to enable finer grain, application specific peering beyond what BGP is capable of today. Since then, various forms and applications of SDX have been proposed. In [24], SDX was defined as a software defined networking exchange that supports interconnect of multiple network domains at flexible layers (not limited to layer 3 or BGP) via signaling among federated network controllers. At the 2013 SDN Program Review workshop [12], one more form of SDX was discussed to be an exchange point for software defined infrastructure (SDI) facilities, providing both SDN interconnect and in-line compute services that can be custom programmed via software APIs.

3 The TARN Architecture

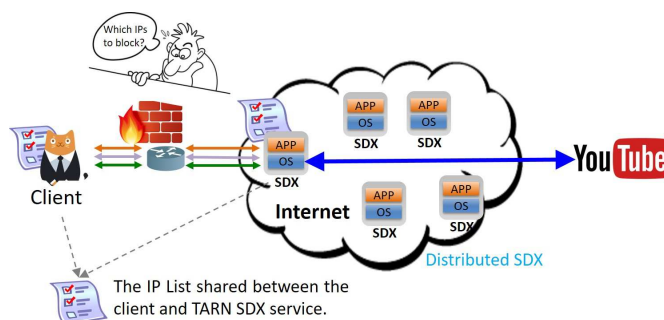


Figure 1: The architecture of TARN for the use case of network censorship circumvention. A client wants to connect to YouTube that is blocked by the client’s local ISP.

3.1 Overview

The core concept of TARN is to bind the two ends of a communication with dynamic, ephemeral and perpetually changing IP addresses. Only authorized users who know the functions used for finding the address “hopping” pattern can track the IP changes over time. This is similar to

frequency-hopping spread spectrum (FHSS) [28] for wireless communications. In this context, we are proposing IP hopping in address space. By doing so, TARN effectively hides the client or/and server from being detected by malicious attackers.

To facilitate an introduction, we consider a use case of TARN – network censorship circumvention. The corresponding TARN architecture is illustrated in Figure 1, in which the client wants to access YouTube which is blocked by the client’s local Internet service provider (ISP). The objective is to enable the client to access YouTube without the ISP detecting and blocking the communication between them. We assume the ISP can monitor all traffic leaving and entering the AS.

In this particular scenario, it is impractical to assume a client would use randomized IP addresses. Most ISPs control the assignment of a fixed IP address to a client. Only the server side (i.e., YouTube) adopts IP hopping in this case. As illustrated in Figure 1, YouTube employs distributed SDX TARN service across the Internet. With the service, YouTube can customize its IP hopping pattern and communicate it with the client as an authorized customer. The SDX provider would provide YouTube a number of different IP prefixes where IP addresses can be flexibly and dynamically used by the YouTube server over time. In this case, the SDX provider will own sufficiently many IP prefixes that can be allocated, most possibly through a lease, to a service provider as YouTube. Note that TARN does not suffer from IP blocking as most other privacy-preserving technologies, including Tor.

3.2 Security features of TARN

The fixed binding between hosts and IP addresses are crucial to today’s Internet routing. However, such scheme has also been taken advantage of by malicious attackers. It enables IP/DNS filtering to become a common method used by nation-state firewalls to block opposition websites. Network users who struggle against censorship really have only two options: VPN and proxy, neither of which can avoid being blocked once being identified. Moreover, fixed IP addresses are also the prime references used by attackers to target their victims, which is a necessary step before staging any malicious attacks.

As a result, breaking the rule of fixed IP addresses becomes a natural solution to thwart network attacks. As the most innovative element of TARN, the idea of IP hopping is similar, but in many aspects superior to the *fast flux* technique used by the notorious botnets to elude detection [16]. And the effectiveness of fast flux as an anti-detection measure has been thoroughly tested and validated in practice by the protracted game between botnets and law enforcement [14]. Therefore, we envision that the adoption of

TARN will mitigate a significant portion of network attacks.

3.2.1 Massaging the IP hopping patterns

We are fully aware that abnormally frequent change of IP addresses may be exploited by adversaries to identify TARN users. We must account for the possible side-channel analysis on the IP hopping frequency of TARN users. In this work, we “massage” the IP hopping pattern of TARN users’ traffic to make it statistically identical to that of the regular network users [34]. This allows traffic generated by TARN users to be seamlessly blended into the background traffic.

The idea is to construct a hidden Markov model (HMM) that represents the IP hopping pattern of regular users and then use the model to determine the dwell time of each IP address for TARN users. The inferred model is known as a deterministic HMM [31, 30, 23]. Unlike the standard HMM that has two sets of random processes: states and observations. In a deterministic HMM, there is only a single set of random processes (time intervals between two hops in this case). So the state transition labels are uniquely associated with an output alphabet.

Given the inferred model, TARN service starts the timing side-channel message by randomly selecting a start state in the model. To determine the dwell time of the current address, a transition is taken from the current state and the corresponding time interval is waited before switching to the next address. If there are more than one possible transitions out of current state, the transition is chosen randomly, weighed on the probability of each transition.

Deep packet inspection (DPI) techniques censor traffic based on the application-layer content and can be used to identify TARN users. We have developed pluggable transports for Tor that mimic protocols like NTP, video games, and smart grid sensors [34]. Our pluggable transports encrypt a session and then translate the AES encrypted data into the syntax of another host protocol. Integration of our pluggable transports into TARN for DPI circumvention is beyond the scope of this paper.

3.3 Implementation Strategies

The TARN architecture can assume multiple forms, from a “host-only” end-to-end solution to a distributed SDX solution. We discuss three different implementation strategies that serve different needs and scopes. All designs are compatible with existing network infrastructure and routing protocols. While implementations vary, they all adopt SDN for programmatic packet switching based on packet flow attributes at different protocol layers (layer 2–layer 4).

Recall that the two core functions of TARN are 1) to unbind IP prefixes from ASes; and 2) to allow sustainable communication sessions between two ends with short-lived

and seemingly randomized IP addresses (not constrained to a single prefix). To achieve this, three design patterns are possible: end-host TARN agent, campus TARN gateway, and SDX TARN service, each of which assumes a different threat model. Table 1 provides a synopsis of the trade-offs for them. Each design pattern implements IP rewriting where appropriate to the specific threat model, i.e., the trust boundary.

3.3.1 End-host TARN agent

In the first pattern, a software agent sits on each end host, together with a virtual SDN switch such as Open vSwitch (OVS) [17], to rewrite the destination and/or source IP addresses before the packets leave the host. The pattern assumes any part of the network between two end hosts are potentially vulnerable to traffic analysis attacks, and thus, requires address rewrite to protect both ends' identity and accessibility. This design pattern typically suffices in allowing a person/client to access a website, e.g., google.com, that would otherwise be blocked by the person's ISP.

Design Pattern	Host application required	BGP route announcements	Trust boundary
End-host TARN agent	✓	✓	Host-only
Campus TARN gateway	✗	✓	Campus network
SDX-based TARN service	✗	✗	Campus network + Path to SDX facilities

Table 1: Considerations for three TARN design patterns.

The TARN agent running on the client-host will rewrite the destination address of the outgoing requests to the website. On the website servers end, either a TARN agent or a TARN gateway (see section 3.3.2) would be used to accept the incoming requests and rewrite them into the correct, original web server addresses. One critical requirement to note – the server side needs to announce via BGP new prefixes that are not the organization's known prefixes. Such prefixes can be scavenged from unused prefixes (a lot more in IPv6 than in IPv4) in the near term. In the long term, we believe SDXs can acquire large pools of IP prefixes for its customers' scheduled reuse. Such prefixes will only be held by any organization for a relatively short time to prevent adversaries from detecting their relationship. The prefixes will only be made known to registered customers through authenticated means to avoid adversaries to fish for the addressing scheme.

3.3.2 Campus TARN gateway

The second pattern relieves the end hosts' burden and runs the IP rewrite actions at a campus gateway router, called the TARN gateway. End hosts choosing to use the service set the TARN gateway their default router. This pattern assumes the entire campus network as a trusted domain where TARN protection is not necessary. This pattern makes more sense to be adopted by major, multi-campus corporates that need to have discrete anonymized inter-campus communications. As a result, the TARN gateways are expected to be deployed at all participating campuses and their functions will be symmetric for any inter-campus communication sessions.

3.3.3 SDX TARN service

As we explore different implementation strategies for TARN, the concept of a distributed SDX surfaces as an attractive option. Recent studies have explored SDX in different forms and for different purposes [18].

Thus far, a SDX has been considered as a SDN-empowered IXP facility - a place where networks of multiple ASes interconnect and forward packet traffic. Today, there are hundreds of IXPs worldwide, and the large IXPs can be interconnecting up to several hundred ASes and forwarding a quarter billion IP addresses each week [11]. In our study, we consider a distributed SDX with the following properties:

- Consisting of multiple distributed facilities spread across Internet;
- Connecting ASes at flexible protocol layers, e.g., as layer-2 circuits or layer-3 routes;
- Supporting custom *packet transformation* in addition to *packet switching*, programmable across any affiliated facilities; and
- Serving each AS as a customer, and facilitating programmable customer-to-customer packet exchange (i.e., transformation and switching).

To achieve these properties, we expect each SDX facility to have a SDN-based switching fabric for interconnecting customer ASes and exchanging traffic between them, and a cloud-like compute infrastructure that executes network functions in between any ASes for custom packet transformation demands. A SDX provider is expected to operate sufficiently many facilities at strategic locations across Internet. Customers for such a SDX service will, most likely through a provider portal, provision a range of network functions at any or all of the provider's facilities to handle incoming and outgoing traffic flows among theirs and others ASes. Figure 2 illustrates the distributed SDX. Gupta




Figure 2: Distributed SDX, where each participant AS connects to a SDX.

et. al have shown in [18] that integrating SDX with existing IXP infrastructure and conventional BGP-speaking ASes is straightforward. And our TARN service makes a compelling wide-area traffic delivery application of SDX.

The SDX pattern consolidates all TARN packet transformation and switching functions to a SDX facility, so the customers are primarily networks, groups of networks or ASes (AS 43515, which is YouTube’s AS number) rather than individual end users. In the case of censorship circumvention, the SDX TARN service appeals to customers like Google, Youtube, NYT, Facebook, Twitter, etc, which are blocked by many nation-state firewalls. From each customer network, be it a campus network or a CDN, we assume a secure channel between the customer and the SDX facility. For example, a layer-2 connection such as the Internet2 Advanced Layer 2 Service (AL2S) [20] can be used. In this case, each customer will have a dedicated landing place which most possibly can be a virtual machine (VM), on which the TARN configurations can be controlled completely by the end-user organization. As a result, the SDX facility can be viewed as a public cloud, with many VMs (at least one per customer) interconnected by the SDN fabric.

3.4 Challenges

We fully recognize that this transformative paradigm introduces significant changes to multiple aspects of the Internet, ranging from routing to system scalability. Since TARN builds on existing Internet infrastructure, to realize TARN in practice, we must identify these challenges and provide sound solutions.

3.4.1 IP address collision avoidance.

TARN works with both IPv4 and IPv6 networks. Address collisions can occur, while the vast IPv6 address space makes this mathematically impossible. Therefore, the approach is particularly flexible with IPv6. The probability of

having at least one collision when randomly mapping the entire IPv4 space into IPv6 space is extremely small, approximately 3.906×10^{-28} compared with the probability of being hit by a meteor being 5.49×10^{-15} .

As discussed later, with SDX, the impact of address collisions can be minimized. Without SDX, collisions will have to be resolved with known duplicate address detection (DAD) methods and packet re-transmissions.

3.4.2 Efficient routing to/between dynamic addresses:

Traffic to these “ephemeral” addresses are routed to the desired destination using BGP route injection, intermediate software defined Internet exchanges (SDX), or a combination of the two. We inject BGP routes for these random and dynamic addresses beforehand and withdraw them after use. In today’s Internet, BGP route updates converge globally in just a few minutes. And the propagation of newly announced prefixes happens almost instantaneously. We are fully aware that ISPs today will not accept frequent BGP announcement, although such operation is technically feasible and compatible with existing Internet infrastructure. The purpose of this work is to demonstrate the technical feasibility and usefulness of this novel approach to pave the way for potential, future global change.

3.4.3 IP address synchronization

IP synchronization between the two ends of an established session is necessary. We use a pseudo-random number generator with a chosen seed to generate a different set of N addresses. To reduce the number of BGP route injections, we limit the generated addresses to be within a smaller range of randomly generated, or deliberately chosen ASes.

The random number generator and seed are and only a properly authenticated client can access via a covert reverse DNS channel [15]. The covert protocol allows us to transform strings into legitimate DNS records. The server encodes the message into a list of domain names and register them to a randomly chosen IP address. The client does a reverse-DNS lookup on the IP address and decodes the domain names to retrieve the message. Different from DNS tunneling, this does not use uncommon record types (e.g., TXT records) or carry suspiciously large volume of traffic as DNS payloads. On the contrary, the resulting traffic will be normal DNS lookup/reverse-lookup traffic, which will not attract attention. The data transmission is not vulnerable to DPI.

4 Experimentation

A series of experiments were conducted to demonstrate the viability and networking capability of this innovative ad-

addressing paradigm. The goal is to show that we can effectively maintain the communication session, while the end host adopts short-lived and randomized IP addresses.

4.1 Testbeds

Our experiment is set up using the networking resources provided by GENI [13] and the PEERING testbed [3]. GENI is a NSF-founded large-scale distributed testbed that gives academics and corporations access to various physical and virtual network resources. With GENI, we can test communications between hosts in different geographic locations across the world. The PEERING testbed is a real BGP testbed that connects to real networks to physical university networks and IXPs around the world. Traditional inter-domain routing experiments were either based on passive observation of existing routes or on simulations, which greatly restricts the fidelity of the results. PEERING removes these limitations by providing researchers a platform to conduct inter-domain routing tests in realistic environments.

4.2 Experiment design

The experiment consists of two parts: BGP route propagation and maintaining the communication session while end host's IP address changes. The configuration of the experiment is shown in Figure 3, where the client and server sit on GENI racks located at NPS and NYSERNet, respectively. On each rack, we installed a virtual network emulator called miniNEXt [27], which extends mininet [2] to easily build complex networks. Each miniNEXt includes a *Floodlight manager* [4] that rewrites packet headers, an *OVS* that interacts with the Floodlight, and a *Quagga router* [6] that runs a virtual BGP router to communicate with the external BGP routers of other ASes.

4.2.1 BGP route propagation

The PEERING testbed maintains multiple IP prefixes for experimental purposes. Researchers can announce BGP prefixes at selected US and EU point of presences (PoPs) to the Internet. As illustrated in Figure 3, the Quagga router at each end is attached to the upstream PEERING server via PEERING's OpenVPN tunnel. PEERING runs its own Quagga software routers to establish a BGP session with their neighbors. Instead of running the BGP route selection process, those PEERING servers provide PEERING clients with full control over BGP route announcements. Our Quagga routers build a BGP peer session with PEERING's Quagga routers. Any BGP announcement from our Quagga router can propagate across the Internet through the PEERING servers, and vice versa. The PEERING server

will just relay the announcements as long as the BGP announcement is confirmed valid by PEERING.

Four IP prefixes were reserved for our experiment: 184.164.243.0/24 for server side and 184.164.242.0/24 for client side. This provides two blocks of contiguous routable IP addresses.

4.2.2 Maintaining communication sessions

In this experiment, only the server changes its address, and the client needs to keep track the server's IP address. This conforms to the use case of censorship circumvention presented in section 3. We leveraged a Floodlight OpenFlow controller that sits on each end to maintain the communication between client and server.

Each Floodlight controller keeps an ordered list of IP addresses called *external IP addresses* that the server is going to use within a certain time frame. In our experiment, a fixed *internal IP* is adopted, which the client/server process talks to directly. This presents the illusion that the application is talking to the same IP all the time while the external IP changes. Floodlight manages the flow rules that rewrite the packet header for port forwarding. The Floodlight controller generates two types of packet: IP packets and ARP packets. With each packet type, there are a pair of flows: one flow rewrites the outgoing packets by changing the internal IP address to the current external IP address; while the other flow rewrites the incoming packets and changes the external IP address to internal IP address. In addition, Floodlight controller is also responsible for updating the external IP address at a pre-defined rate that is only known to the client.

4.3 Experimental Results

Since TARN is an ongoing project, only preliminary experimental results are provided. It is assumed that a pseudorandom sequence of IP addresses are known to both the server and client. It is also assumed that the server rapidly changes its IP address in a predetermined order that is known to both ends. We have provided methods of IP synchronization between the client and server in section 3.4. The corresponding module will be implemented and integrated in future work. The configuration of the experiment is consistent with the censorship circumvention use case in Figure 1, where the server's IP address changes at a fixed rate.

The objective of this experiment is to demonstrate two basic functions of TARN: (1) The SDN controller offers the capability to manipulate the hopping pattern of IP addresses, and (2) TARN maintains a stable and fluent session while the server frequently changes its IP address (111 different IP addresses from 184.164.243.0/24 with average dwell time between hops < 10 sec).




Figure 3: Configuration of the experiment.

We had the client send out 672 IP packets and collected data received on the server. All the packets were successfully received by the client. Given that the packets were sent at a stable rate, this means the dwell time for each IP address the server used were randomized. Moreover, any given distribution of IP hopping can be achieved using an SDN controller. Although only preliminary results are available, they fully support that we can effectively maintain the communication session, while the end host adopts short-lived and randomized IP addresses. This lays the foundation for further experimentation, including hopping pattern manipulation and two-end IP hopping.

5 Conclusion and Future Work

The TARN system is highly resistant DNS/IP filtering and traffic analysis, by removing a major vulnerability in today’s addressing scheme – the fixed binding between hosts and IP addresses, which has been used by cyber attacks to easily locate their victims. TARN relies on the promising SDN technologies to realize stable and efficient routing between end hosts with dynamic addresses.

We lay out the technical foundation and roadmap for a scalable implementation of this transformative paradigm. Three different implementation strategies are provided, each of which performs IP hopping where appropriate to the users’ specific requirements given the trust boundary. We show that TARN is compatible with today’s Internet infrastructure by properly addressing the possible challenges in deploying TARN.

The one-way randomization version of the SDN-based solution was implemented and tested on PEERING’s BGP testbed. Since TARN is an ongoing project, only preliminary experimental results are provided in this paper. The obtained results demonstrate that: (1) SDN controllers allow users to easily manipulate the IP hopping pattern to their own needs; and (2) it is completely feasible to maintain a stable and fluent communication between end hosts of dynamic IPs. This has laid the foundation for further experimentation.

Future work includes (1) performing experiments to demonstrate the utility of the other designs, (2) completing the two-way IP randomization, (3) performing an in-depth security analysis of the SDN-based TARN solution, and (4) implementing TARN as the future of SDX. We will perform extensive, invasive testing to rigorously prove that TARN is to IP/DNS filtering, DPI and side-channel attacks. These results, in conjunction with our preliminary security analysis, show great strides towards a security-centered Internet without compromising performance.

Acknowledgment

This material is based upon work sponsored by the National Science Foundation under Grant No. 1643020. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

- [1] The invisible internet project. <https://psiphon.ca/en/download.html>.
- [2] Mininet: An instant virtual network on your laptop (or other pc). <http://mininet.org/>.
- [3] Peering – the bgp testbed. <https://peering.usc.edu/>.
- [4] Project floodlight. <http://www.projectfloodlight.org/>.
- [5] Psiphon. <https://psiphon.ca/en/download.html>.
- [6] Quagga routing suite. <http://www.nongnu.org/quagga/>.
- [7] Tor. <https://www.torproject.org/projects/torbrowser.html.en>.
- [8] Freedom in the world 2017, 2017. <https://freedomhouse.org/report/freedom-world/freedom-world-2017>.
- [9] C. Beasley, X. Zhong, J. Deng, R. Brooks, and G. Kumar Venayagamoorthy. A Survey of Electric Power Synchronphasor Network Cyber Security. In *Innovative Smart Grid Technologies Conference Europe (ISGT-Europe), 2014 IEEE PES*, pages 1–5, Oct 2014.
- [10] R. Brandom. A network error routed traffic for the uk’s nuclear weapons agency through russian telecom, March 2015. <https://www.theverge.com/2015/3/13/8208413/uk-nuclear-weapons-russia-traffic-redirect>.

- [11] N. Chatzis, G. Smaragdakis, J. Böttger, T. Krenc, and A. Feldmann. On the benefits of using a large ixp as an internet vantage point. In *Proceedings of the 2013 conference on Internet measurement conference*, pages 333–346. ACM, 2013.
- [12] V. Dattoria, I. Monga, B. Lyles, K. Thompson, and G. Miller. Sdn program review final report, December 2013. <https://www.orau.gov/sdnpr2013/>.
- [13] N. S. Foundation. Geni (global environment for network innovations). <http://www.geni.net/>.
- [14] Y. Fu, B. Husain, and R. R. Brooks. Analysis of botnet counter-counter-measures. In *Proceedings of the 10th Annual Cyber and Information Security Research Conference*, page 9. ACM, 2015.
- [15] Y. Fu, Z. Jia, L. Yu, X. Zhong, and R. Brooks. A covert data transport protocol. In *Malicious and Unwanted Software (MALWARE), 2016 11th International Conference on*, pages 1–8. IEEE, 2016.
- [16] Y. Fu, L. Yu, O. Hambolu, I. Ozcelik, B. Husain, J. Sun, K. Sapra, D. Du, C. T. Beasley, and R. R. Brooks. Stealthy domain generation algorithms. *IEEE Transactions on Information Forensics and Security*, 12(6):1430–1443, 2017.
- [17] L. Foundation. Open vswitch. <http://openvswitch.org/>.
- [18] A. Gupta, L. Vanbever, M. Shahbaz, S. P. Donovan, B. Schlinker, N. Feamster, J. Rexford, S. Shenker, R. Clark, and E. Katz-Bassett. Sdx: A software defined internet exchange. *ACM SIGCOMM Computer Communication Review*, 44(4):551–562, 2015.
- [19] A. Houmansadr, T. J. Riedl, N. Borisov, and A. C. Singer. I want my voice to be heard: Ip over voice-over-ip for unobservable censorship circumvention. In *NDSS*, 2013.
- [20] Internet2. Layer 2 services. <https://www.internet2.edu/products-services/advanced-networking/layer-2-services/>.
- [21] P. Kampanakis, H. Perros, and T. Beyene. Sdn-based solutions for moving target defense network protection. In *World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2014 IEEE 15th International Symposium on a*, pages 1–6. IEEE, 2014.
- [22] J. Karlin, D. Ellard, A. W. Jackson, C. E. Jones, G. Lauer, D. Mankins, and W. T. Strayer. Decoy routing: Toward unblockable internet communication. In *FOCI*, 2011.
- [23] C. Lu, J. M. Schwier, R. M. Craven, L. Yu, R. R. Brooks, and C. Griffin. A normalized statistical metric space for hidden markov models. *IEEE transactions on cybernetics*, 43(3):806–819, 2013.
- [24] J. Mambretti, J. Chen, and F. Yeh. Software-defined network exchanges (sdxs): Architecture, services, capabilities, and foundation technologies. In *Teletraffic Congress (ITC), 2014 26th International*, pages 1–6. IEEE, 2014.
- [25] I. Ozcelik, Y. Fu, and R. R. Brooks. Dos detection is easier now. In *Research and Educational Experiment Workshop (REE), 2013 Second GENI*, pages 50–55. IEEE, 2013.
- [26] H. Roberts, E. Zuckerman, and J. Palfrey. Circumvention landscape report: Methods, uses, and tools. *Retrieved August*, 5:2012, 2007.
- [27] B. Schlinker. mininext. <https://github.com/USC-NSL/miniNExT>.
- [28] S. M. Schwartz. Frequency hopping spread spectrum (fhss) vs. direct sequence spread spectrum (dsss) in the broadband wireless access and wlan arenas. *white paper*, 2001.
- [29] C. Tang. In-depth analysis of the great firewall of china, 2016. <http://www.cs.tufts.edu/comp/116/archive/fall2016/ctang.pdf>.
- [30] L. Yu. Stochastic tools for network security: Anonymity protocol analysis and network intrusion detection. 2012.
- [31] L. Yu, J. M. Schwier, R. M. Craven, R. R. Brooks, and C. Griffin. Inferring statistically significant hidden markov models. *IEEE Transactions on Knowledge and Data Engineering*, 25(7):1548–1558, 2013.
- [32] K. Zeeter. Someone’s been siphoning data through a huge security hole in the internet, April 2017. <https://www.wired.com/2013/12/bgp-hijacking-belarus-iceland/>.
- [33] X. Zhong, A. Ahmadi, R. Brooks, G. K. Venayagamoorthy, L. Yu, and Y. Fu. Side channel analysis of multiple pmu data in electric power systems. In *Power Systems Conference (PSC), 2015 Clemson University*, pages 1–6. IEEE, 2015.
- [34] X. Zhong, Y. Fu, L. Yu, R. Brooks, and G. K. Venayagamoorthy. Stealthy malware traffic-not as innocent as it looks. In *Malicious and Unwanted Software (MALWARE), 2015 10th International Conference on*, pages 110–116. IEEE, 2015.
- [35] X. Zhong, I. Jayawardene, G. K. Venayagamoorthy, and R. R. Brooks. Denial of service attack on tie-line bias control in a power system with pv plant. *IEEE Transactions on Emerging Topics in Computational Intelligence*.