Global Internet Traffic Routing and Privacy

Richard Brooks
Holcomb Dept. of Electrical and Computer Engineering
Clemson University
Clemson, SC, USA
rrb@clemson.edu

Kuang-Ching Wang
Holcomb Dept. of Electrical and Computer Engineering
Clemson University
Clemson, SC, USA
kwang@clemson.edu

Jon Oakley
Holcomb Dept. of Electrical and Computer Engineering
Clemson University
Clemson, SC, USA
joakley@clemson.edu

Nathan Tusing
Holcomb Dept. of Electrical and Computer Engineering
Clemson University
Clemson, SC, USA
ntusing@clemson.edu

***Abstract*—Current Internet Protocol routing provides minimal privacy, which enables multiple exploits. The main issue is that the source and destination addresses of all packets appear in plain text. This enables numerous attacks, including surveillance, man-in-the-middle (MITM), and denial of service (DoS). The talk explains how these attacks work in the current network. Endpoints often believe that use of Network Address Translation (NAT), and Dynamic Host Configuration Protocol (DHCP) can minimize the loss of privacy. We will explain how the regularity of human behavior can be used to overcome these countermeasures. Once packets leave the local autonomous system (AS), they are routed through the network by the Border Gateway Protocol (BGP). The talk will discuss the unreliability of BGP and current attacks on the routing protocol. This will include an introduction to BGP injects and the PEERING testbed for BGP experimentation. One experiment we have performed uses statistical methods (CUSUM and F-test) to detect BGP injection events. We describe work we performed that applies BGP injects to Internet Protocol (IP) address randomization to replace fixed IP addresses in headers with randomized addresses. We explain the similarities and differences of this approach with virtual private networks (VPNs). Analysis of this work shows that BGP reliance on autonomous system (AS) numbers removes privacy from the concept, even though it would disable the current generation of MITM and DoS attacks. We end by presenting a compromise approach that creates software-defined data exchanges (SDX), which mix traffic randomization with VPN concepts. We contrast this approach with the Tor overlay network and provide some performance data.**

***Index Terms*—BGP Injection, SDN, MitM, Privacy**

I. Introduction

Today's Internet is a global infrastructure that supports finance, business, research, politics, journalism, entertainment, and private personal communications. These applications are subject to surveillance, filtering, and tampering by attackers anywhere on Earth. Attackers can be (and are) criminals, governments, jealous partners, voyeurs, business competitors, private companies, . . . , and combinations thereof. Routing insecurity enables:

- Countries and corporations routinely perform Domain Name System (DNS) [12], [21] and IP address [10] black-list filtering to block users from accessing network address ranges.
- Deep packet inspection (DPI) [4], [26] blocks network streams, sometimes by inserting reset packets when keywords of interest are detected [8], [24].
- Denial of Service (DoS) attacks deny all legitimate access to a service [20].
- Man in the Middle attacks occur by intercepting network connections and placing malicious logic in the middle of a network connection [9].

The basic problem is that network connections are treated as non-sensitive information. Little is done to keep this information confidential. In the current Internet architecture:

- DNS is a global database mapping clear-text symbolic node names to clear-text IP addresses. Most DNS traffic is currently in clear text; subject to surveillance; and vulnerable to man-in-the-middle attacks. Even when the "last mile" is encrypted, users access DNS request information from local ISPs or open DNS servers with their own legal restrictions. ISP's are subject to local regulations. Note that for companies, DNS traffic logs reveal sensitive information about internal R&D efforts.
- Communications, except for traffic using tools like Tor and Psiphon, are routed directly from one node to another using source and destination IP address information that is available in clear-text in the packet header. IP addresses belong to specific entities. It is trivial to block access to sites, like the New York Times, by blocking all traffic to and from the New York Times range of IP addresses.[1]
- Traffic monitoring and profiling are easy. The set of DNS and IP addresses accessed are easily tracked by anyone with access to the regional network. Classes of communications can either be read directly from IP port numbers.
- Global traffic can be almost arbitrarily rerouted through misuse of the Border Gateway Protocol (BGP). BGP defines IP routing paths by propagating perceived distances between autonomous systems. Packets take the shortest route from packet source to destination, as defined by the hop's BGP tables. In one example of this abuse, China arbitrarily hijacks US Internet traffic steering it into China for later analysis [11].
- denial of service attacks are trivial. Since there is no filtering of data entering the network globally, it is easy to introduce an excessive volume of network traffic aimed at the IP address of a site that is considered undesirable.

This material is based upon work sponsored by the National Science Foundation under Grant No. 1643020.

[1]The existence of content delivery networks can make filtering IP addresses slightly more complicated, but not enough to make traffic filtering and DoS less widespread.

978-1-7281-6460-1/20/$31.00 ©2020 IEEE

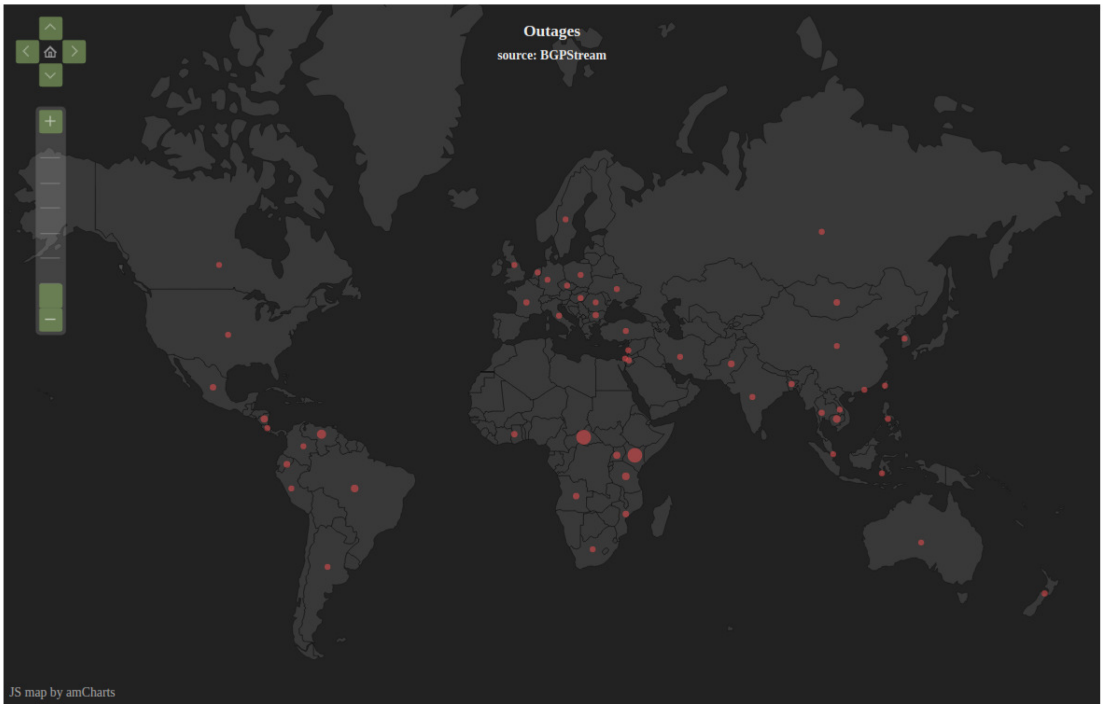

Fig. 1. BGP incidents logged at bgpstream.com on August 28, 2020.

> Root DNS servers have been disabled; credit card companies have suffered an economic loss, and many media sites have been targeted [20].

Some efforts have been attempted to make it difficult to tamper with DNS and BGP infrastructure, but little has been done to make meta-data about connections confidential.

VPNs encrypt all traffic, including IP addresses, but traffic can still be inferred using side-channel attacks (timing, packet size) [15], [28]. As [3], [28] found, these solutions are imperfect. There are also solution specific issues–Lantern is only active when a website is blocked [23], leading to a myriad of potential attacks. In practice, VPN companies must choose between turning over logs or facing federal charges [25]. In all of these cases, the users' privacy is in the hands of their chosen solution. Additionally, these solutions are easily detected with IP blacklist or PCAP-based rules to detect VPNs.

Since proxies and VPNs fail to provide sufficient privacy in several cases, anonymity networks like Tor [2] and I2P [1] have arisen. Tor's Onion Routing encrypts traffic at least three times, letting only the current node know the next destination. I2P is not widely used, despite being similar to Tor in many ways. There have been proof-of-concept attacks against the anonymity of Tor users.

This paper discusses vulnerabilities with establishing and maintaining Internet connections and innovative ways to remove the flaws that make them vulnerable while maintaining existing infrastructure.

II. BGP HIJACKING

Frequent errors occur with BGP routing. Figure 1 is a map showing the locations of BGP errors detected around August 28, 2020. BGP allows AS's to advertise IP prefixes they serve and routes its neighbors can reach through them. Major telecom service providers are labeled as *tier 1*. They have a direct connection to the Internet backbone. IP traffic is served according to a peer-or-pay system, where networks either provide services for each other (peer) or have to pay when their traffic moves through another network. Tier 1 nodes reciprocally share massive volumes of data without paying fees [7].

BGP trusts AS's to only advertise IP ranges they own and legitimate paths. Unfortunately, trust is not always merited; causing IP traffic to route incorrectly. Four general classes of IP hijacking are [7]:

1) typographical errors;
2) prepending mistakes;
3) origin changes; and
4) forged AS paths.

The first two are typically due to human error, also known as *fat fingering*. The last two are more often malicious. Figure 2 is a graph showing information related to a Russian AS's BGP leak.

For example, China Telecom has ten Internet *points of presence* (PoPs) located on the Internet backbone in North America. Eight are in the USA and two in Canada. The USA has no PoPs in China. This allows China to route North American IP traffic into its network at will. Such as [11]:

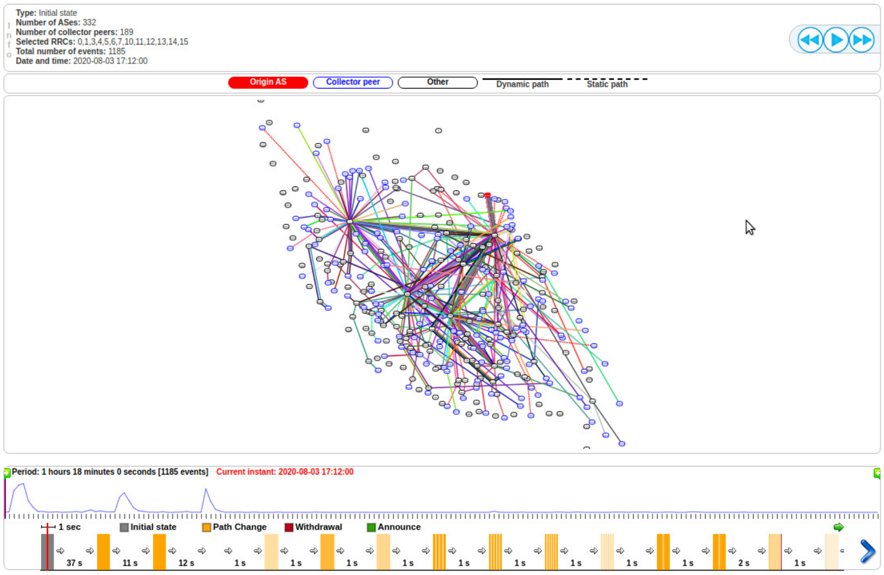

Fig. 2. Example BGP incident display for Russian AS.

- April 8, 2010 china Telecom rerouted 15% of IP traffic into China for 18 minutes, which was believed to be a proof-of-concept test that it can reroute traffic at will;
- Starting in February 2016, Internet traffic between the Canadian and South Korean governments were routed through the Chinese mainland for six months. This detour has re-appeared since then for shorter duration;
- In October 2016, traffic from several USA financial IP sites to their European headquarters in Milan was redirected through China. It appears that no route was found from China to Milan; leading the traffic to be eventually discarded.
- In April and May of 2017, Internet traffic for a US news organization was routed through China from Scandinavia to Japan.

The supposition is that these traffic detours were made to enable either traffic analysis or MITM. Other notable BGP misdirection incidents include:

- In 2008, Pakistan sink-holed YouTube for the entire world in an effort to make content contrary to Islam [11]; unreachable [11];
- In 2018, Google traffic was routed through Nigeria, China, and Russia [27]; and
- An August 28, 2020 review of [5] reports possible hijacks in the USA, Brazil, Poland, UK, Portugal, Latvia, Argentina, Peru, Australia, Singapore,. . . . With suspected attack sources being just as varied.

Modifications to BGP exist that use public-key cryptography to associate IP address ranges with AS's, but as of September 2020, a current list of 28 major operators [2] only shows 8 of them fully participating. There are multiple reasons for slow adoption of these BGP modifications [13]. These include costs and the fact that BGP update rules include informal social agreements.

BGP routing insecurity means that IP traffic can be routed arbitrarily. This includes possible denial of service for any IP address from any location by sinkholing, as illustrated by Pakistan's inadvertent sinkholing of YouTube for the whole world. It allows web traffic to/from any specific IP range to be observed and recorded anywhere on Earth. Man-in-the-middle attacks are enabled by intentionally routing IP traffic into network segments controlled by malicious parties.

To date, BGP routing misbehavior is detected by historical analysis of BGP updates. The PEERING testbed[3] provides access to the BGP routing infrastructure for research. We used PEERING to simulate BGP hijacking for active TCP sessions. The endpoints for the sessions remained constant, but BGP updates would modify the traffic paths.

We hypothesized that modifying the paths taken by the packets would likely change traffic dynamics. We used PEERING to test this hypothesis. To avoid abuse, PEERING requires a strict definition of the experiments. We were given access to three points-of-presence. The experiment set up TCP connections that mimicked traffic from Phasor Measurement Units. During the sessions, we used BGP injects to reroute the TCP sessions. The route between the source and destination AS's moved to another continent.

Changing the route taken by a TCP session modifies the routers that forward the traffic, the number of routers in the path, and the interactions between various Internet data streams. We measured the following attributes that should reflect these changes:

1) Path latency;
2) Path latency variance (Jitter);

[2]https://isbgpsafeyet.com/

[3]https://peering.usc.edu

3) Inter-packet delay; and
4) Inter-packet delay variance (jitter).

To detect these changes, we use CUSUM change point detection and the F-Test from statistics. In [6], [19], CUSUM detects Distributed Denial of Service (DDoS) attacks. The F-Test is an established statistical test to see if two samples have the same variance.

CUSUM detects significant mean value changes hidden in noise. Detection uses a sequential probability ratio test (SPRT). The modified CUSUM algorithm is:

$$\tilde{S}[t] = max\{0, (\tilde{S}[t-1] + |\tilde{m}^S[t] - m^L[t]| - C)\}; \quad \tilde{S}[0] = 0$$

$S[t-1]$ is the old CUSUM value, $m^S[t]$ is the short window average of packets' latency, $m^L[t]$ is the long-term average of packets' latency, calculated with a given long term average memory parameter ε, $0 < \varepsilon < 1$:

$$m^L[t] = \varepsilon m^L[t-1] + (1-\varepsilon)m^S[t]; \quad m^L[0] = 0$$

To reduce high frequency noise, local averaging uses α to create a low-pass filter:

$$\tilde{m}^S[t] = \alpha m^S[t] + (1-\alpha)\tilde{m}^S[t-1]; \quad \tilde{m}^S[0] = 0$$

C is a correction parameter that forces small CUSUM values to 0. $\tilde{S}[t]$ will increase when the short term average is consistently significantly larger than the long term average.

The F-Test is an established statistical hypothesis test to see if two samples have the same variance [17]. We check to see if the current time series variance is the same as the time series's historical variance. Let $v^S[t]$ be the variance of the current data sample with n_s samples at time t. Let $v^L[t]$ be the historical variance of the data sample with $n_L > n_S$ samples at time t. Then the F-Test statistic is:

$$(v^L[t])^2/(v^S[t])^2$$

The critical value found using a table of F statistics for $n_L - 1, n_S - 1$ degrees of freedom with 95% confidence, the hypothesis that two data set have the same variance can be rejected or accepted.

We apply CUSUM and F-Test to absolute and inter-packet delays of the captured traffic. ROC curves are shown in Figs 3a to 3d.

The ROC curves indicate that CUSUM detects BGP route change using mean traffic latency with high True Positive Rate (TPR) and low False Positive Rate (FPR). However, as the inter-packet delay mean is not significantly affected by BGP route changes, CUSUM analysis of inter-packet delays are not effective. The F-Test detects the BGP route change on both the inter-packet and absolute delay. Compared to the CUSUM test; the F-Test provides higher TPR and FPR on inter-packet delays. This analysis is based on experimental data.

Unfortunately, the experimentation platform limited us to a small number of points of presence and AS's. While we hypothesize that this approach might allow network users to identify when network connections are subjected to BGP hijacking.

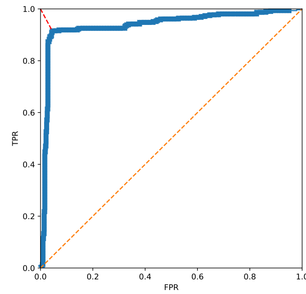

(a) CUSUM ROC for PMU packets absolute delay for BGP hijacking detection. The best operation point gives 91.56% TPR and 4.47% FPR

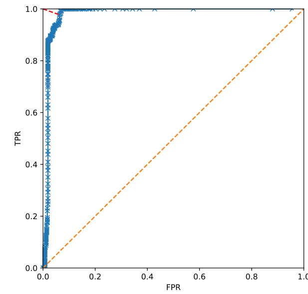

(b) F-Test ROC for PMU packets absolute delay for BGP hijacking detection. The best operation point gives 97.72% TPR and 6.35% FPR

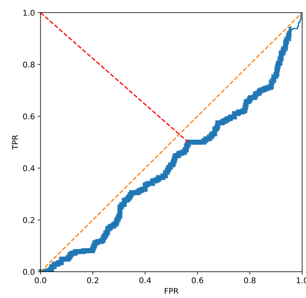

(c) CUSUM ROC for PMU packets inter-packet delay for BGP hijacking detection. The best operation point gives 50% TPR and 56% FPR

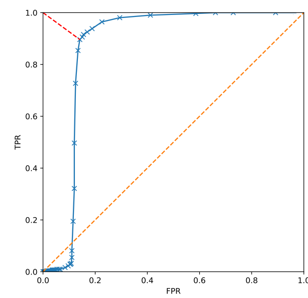

(d) F-Test ROC for PMU packets inter-packet delay for BGP hijacking detection. The best operation point gives 89.61% TPR and 14.12% FPR

III. User Attribution

Murdoch and Danezis used statistical traffic analysis to deanonymize the Tor network by measuring the load on relay nodes [16]. Øverlier and Syverson used timing-based correlations to deanonymize hidden services [18]. Johnson et al. extend these ideas with user behavior and common services to show realistic adversary can deanonymize Tor [14]. We use BGP injection as the basis for more complex traffic analysis. An observer can redirect traffic en-route and use traffic metadata to classify users.

To show user attribution, we designed an experiment to classify traffic from three users. Two characteristics defined the simulated traffic from each user–size and interpacket delay. The size is the full size of the packet in bytes (including header), and the interpacket delay is the time between two packets in seconds. The traffic from each user is combined and sent to an observer that infers the user from traffic characteristics. User i's traffic characteristics are defined in Equation 1.

$$\begin{aligned} \text{User } i \text{ Packet Size} &= N(\mu_{s,i}, \sigma_{size}) \\ \text{User } i \text{ IPD} &= N(\mu_{ipd,i}, \sigma_{ipd}) \end{aligned} \tag{1}$$

We considered two cases, shown in Table I. For the size, σ_{size} was 250 bytes, and for the interpacket delay, σ_{ipd} was 0.01 seconds.

TABLE I
USER CHARACTERISTIC PARAMETERS.

	$1-\sigma$		$3-\sigma$	
User	Size μ_i (kB)	IPD μ_i (s)	Size μ_i (kB)	IPD μ_i (s)
1	1250	0.04	1000	0.03
2	1500	0.05	1750	0.06
3	1750	0.06	2500	0.09

The simulated user traffic for the user characteristics three standard deviations apart is visualized in Figure 3a and Figure 3b. Since the traffic is combined, some information conveyed by the IPD is lost. Figure 3a shows no distinct peaks that represent each user.

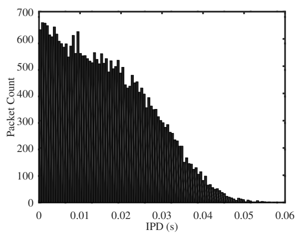

(a) Observed interpacket delay with three users (3-σ separation).

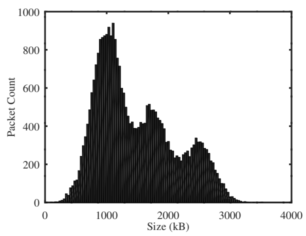

(b) Observed packet size with three users (3-σ separation).

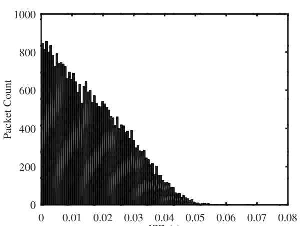

(c) Observed interpacket delay with three users (1-σ separation).

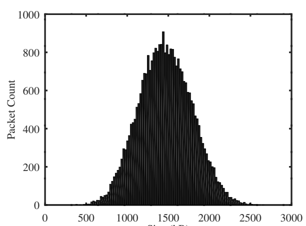

(d) Observed packet size with three users (1-σ separation).

Fig. 3. Two sets of characteristics (size and IPD) used to model user traffic.

In contrast, Figure 3b shows three distinct peaks that represent each user. The first user's peak is higher because of the smaller interpacket delay, which led to more observations. The simulated user traffic for the user characteristics one standard deviations apart is visualized in Figure 3c and Figure 3d. In this case, Figure 3c still shows no distinct peaks, despite the mean IPDs being closer. The observed packet size distribution also has no distinct peaks, as expected in such a closely clustered case.

We used sklearn [22] to test two classifiers against this data: a simplistic *a priori* minimum distance classifier and a random forest classifier (RFC). The higher IPD led to label imbalance, as evident in Figure 3b, so the weighted classifier metrics were used to evaluate performance. We used a 66%-33% split of the train-test data to train the sklearn classifiers. The *a priori* minimum distance classifier is defined in Equation 2.

$$\text{Min Distance}(\theta) = i \quad \text{s.t.} \quad |\mu_i - \theta| \leq |\mu_k - \theta| \quad \forall k \tag{2}$$

TABLE II
CLASSIFIER METRICS USING ONLY PACKET SIZE.

	$1-\sigma$		$3-\sigma$	
Metric	Min. Distance	GNB	Min. Distance	GNB
Precision	60	61	92	92
Recall	59	59	92	92
F_1	60	60	92	92

TABLE III
CLASSIFIER METRICS USING ONLY IPD.

	$1-\sigma$		$3-\sigma$	
Metric	Min. Distance	RFC	Min. Distance	RFC
Precision	34	51	34	72
Recall	34	50	33	58
F_1	34	50	31	62

If each user's characteristics are known to the observer, this classifier chooses the user with the mean closest to the observed value. Table II shows the classification results using the size feature alone. We compared the minimum distance classifier to a Gaussian Naive-Bayes (GNB) classifier, which did not know the user characteristics.

Both classifiers perform well when three standard deviations separate the user characteristics, but performance degrades when only a single standard deviation separates them.

When considering only IPD, classification becomes harder. Table III shows the classification results using the same minimum distance classifier compared to a random forest classifier. We calculate the minimum distance between the average IPD for each user and the cumulative sum of each IPD in a given window. This classifier is defined in Equation 3.

$$\begin{aligned} &\text{Min Distance}(\hat{\boldsymbol{\theta}}) = i \\ &\text{s.t.} \quad |\mu_i - \Sigma_{j=0}^{M}\theta_j| \leq |\mu_k - \Sigma_{j=0}^{M}\theta_j| \quad \forall k \\ &\text{where} \quad \hat{\boldsymbol{\theta}} = \{\theta_0, \theta_1, ..., \theta_M\} \end{aligned} \tag{3}$$

The random forest classifier was successful using the IPD characteristic as a feature. In both the one and three standard deviation separations, the minimum distance classifier was no better than guessing, but the random forest classifier provided promising results.

The minimum distance classifier and the random forest classifier used all features for user classification. The Gaussian Naive-Bayes classifier was used as a meta-classifier to combine the two minimum distance classifications. The random forest classification used both size and IPD features.

TABLE IV
CLASSIFIER METRICS USING IPD AND PACKET SIZE.

	$1-\sigma$		$3-\sigma$	
Metric	Min. Distance w/ GNB	RFC	Min. Distance w/ GNB	RFC
Precision	64	65	92	93
Recall	58	63	92	93
F_1	60	64	92	93

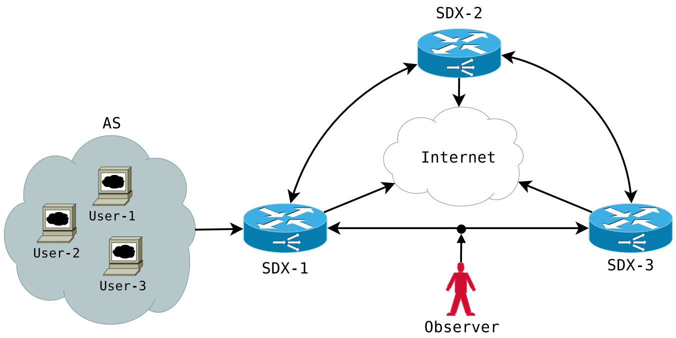

Fig. 4. TARN architecture for providing privacy using BGP injection principals.

Since the was so much information available in the size characteristic, the classifier using both size and IPD tended towards the results obtained when *only* using size as a feature. The classifiers performed similarly in both the one standard deviation case and the three standard deviation cases.

Regardless of the classifier used, it is possible to deanonymize users to a degree using the most simplistic characteristics. The de-anonymization becomes more challenging as the characteristics become closer together (one standard deviation separation versus three standard deviations).

IV. Mitigations

The Traffic Analysis Resistant Network (TARN) is an SDX-based architecture that addresses traffic analysis and this form of user classification, using the principals of BGP injection discussed in Section II. Encrypting traffic provides a degree of privacy, but the classification discussed in Section III work regardless of encryption[4]. TARN has several fundamental properties that counteract these (and many other) de-anonymization techniques using ideas derived from BGP injections. Figure 4 shows the high-level TARN architecture.

User traffic from an AS is sent(via a secure L2 connection) to a TARN SDX node. Each SDX node connects to ASes that it services, an internet gateway, and other TARN SDX nodes. TARN randomly routes user traffic through several other TARN SDX nodes before sending it over the internet. Further, a single flow could get split between multiple TARN nodes. Each TARN node is addressable via other TARN nodes by a set of IP addresses. The set of IP addresses assigned to each node changes on a fixed interval, so traffic between nodes appears to have random IP addresses. Observing BGP announcements would allow an attacker to abstract the TARN node to the AS using the prefix, but eBGP allows ASes to share IP prefixes, and it can obfuscate the AS associated with the source and destination.

TARN encrypts the connection between nodes and uses a fixed packet size. The resulting histogram of traffic sizes for all users resembles Figure 5, and it leaves no information for a classifier (*a priori* or otherwise) to use to classify user traffic. TARN nodes combine user traffic and randomly send it to another TARN node. In contrast, moving target defense changes the perceived network topology dynamically. Thus far, there have been no infrastructure-based moving target defense solutions. All existing solutions rely on controlling the infrastructure. Classifying TARN traffic approximates the case of trying to classify traffic by IPD when the IPDs are very similar (Figure 3c). From Section III, we showed that in the simplest case, the random forest classifier was virtually equivalent to random guessing.

[4]VPN packet sizes are usually rounded up to the nearest block size, which is still an effective classification feature

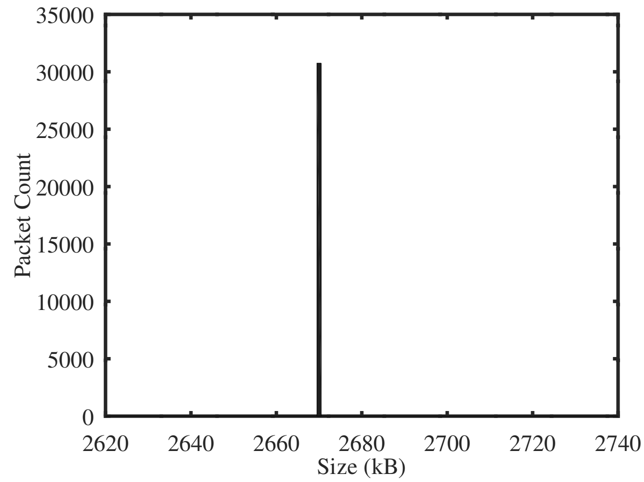

Fig. 5. All user packet sizes in a TARN connection.

V. Conclusions

Current Internet Protocol routing provides minimal privacy, which enables multiple exploits. BGP injection presents a real threat to any user sending traffic over the internet. Regardless of source or destination, BGP injections can allow an attacker to redirect traffic to their infrastructure.

Currently, BGP malfeasance detection is done offline by analyzing historical traces of BGP injection traffic. We tested the hypothesis that online analysis of IP traffic dynamics could provide a reasonable detection metric for BGP tampering. This hypothesis warrants larger-scale testing when experimental resources become available. Ideally, these dynamic features could alert the AS that is being attacked. The AS could examine relevant BGP injection traffic for verification. At which point, technical and social countermeasures could be undertaken in real-time. Further research and testing are needed before this vision can be realized.

Traffic analysis presents a severe threat, even with existing privacy tools. Our user classification experiment used two general traffic characteristics packet size and interpacket delay (IPD). Using both these features, we showed that both an *a priori* classifier and a random forest classifier were effective at classifying user traffic with distinct features (three standard deviation separation). When the traffic was more uniform (one standard deviation of separation), both classifiers' performance deteriorated, but they were still better than randomly guessing.

By classifying using only these features, we show that most privacy tools are vulnerable to this form of classification unless they actively change them to prevent analysis.

Finally, we present an SDX-based Traffic Analysis Resistant Network (TARN) solution that uses BGP injection's fundamental principles to prevent analysis and user classification. Each TARN node uses pseudo-random IP addresses with a dynamic eBGP announcement to hide the traffic's real source and destination. TARN's uniform encrypted packet size and user traffic splitting (effectively randomizing IPD) obscure even the most basic classification features. TARN presents a unique infrastructure-based solution to privacy that leverages one of the most insecure aspects of the modern internet.

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://geti2p.net/en/.

[2] Tor. https://www.torproject.org/.

[3] Noah Apthorpe, Dillon Reisman, Srikanth Sundaresan, Arvind Narayanan, and Nick Feamster. Spying on the smart home: Privacy attacks and defenses on encrypted IoT traffic. *arXiv preprint arXiv:1708.05044*, 2017.

[4] Simurgh Aryan, Homa Aryan, and J Alex Halderman. Internet censorship in iran: A first look. In *3rd USENIX Workshop on Free and Open Communications on the Internet (FOCI 13)*, 2013.

[5] BGPStream. Bgpstream, https://BGPStream.com, (last visited 08/2020), 2020.

[6] Glenn Carl, Richard R. Brooks, and Suresh Rai. Wavelet based denial-of-service detection. *Comput. Secur.*, 25(8):600–615, November 2006.

[7] Shinyoung Cho, Romain Fontugne, Kenjiro Cho, Alberto Dainotti, and Phillipa Gill. Bgp hijacking classification. In *2019 Network Traffic Measurement and Analysis Conference (TMA)*, pages 25–32. IEEE, 2019.

[8] Richard Clayton, Steven J Murdoch, and Robert NM Watson. Ignoring the great firewall of china. In *International Workshop on Privacy Enhancing Technologies*, pages 20–35. Springer, 2006.

[9] Mauro Conti, Nicola Dragoni, and Viktor Lesyk. A survey of man in the middle attacks. *IEEE Communications Surveys & Tutorials*, 18(3):2027–2051, 2016.

[10] Ronald J Deibert. The geopolitics of internet control: Censorship, sovereignty, and cyberspace. *Routledge handbook of Internet politics*, pages 323–336, 2009.

[11] Chris C Demchak and Yuval Shavitt. China's maxim–leave no access point unexploited: The hidden story of china telecom's bgp hijacking. *Military Cyber Affairs*, 3(1):7, 2018.

[12] Andrea Di Florio, Nino Vincenzo Verde, Antonio Villani, Domenico Vitali, and Luigi Vincenzo Mancini. Bypassing censorship: A proven tool against the recent internet censorship in turkey. In *2014 IEEE International Symposium on Software Reliability Engineering Workshops*, pages 389–394. IEEE, 2014.

[13] Sharon Goldberg. Why is it taking so long to secure internet routing? *Communications of the ACM*, 57(10):56–63, 2014.

[14] Aaron Johnson, Chris Wacek, Rob Jansen, Micah Sherr, and Paul Syverson. Users get routed: Traffic correlation on tor by realistic adversaries. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 337–348, 2013.

[15] Hongda Li, Fuqiang Zhang, Lu Yu, Jon Oakley, Hongxin Hu, and Richard R Brooks. Towards efficient traffic monitoring for science dmz with side-channel based traffic winnowing. In *Proceedings of the 2018 ACM International Workshop on Security in Software Defined Networks & Network Function Virtualization*, pages 55–58. ACM, 2018.

[16] Steven J Murdoch and George Danezis. Low-cost traffic analysis of Tor. In *Security and Privacy, 2005 IEEE Symposium on*, pages 183–195. IEEE, 2005.

[17] John Neter, William Wasserman, and Michael H Kutner. Applied linear regression models. 1989.

[18] Lasse Overlier and Paul Syverson. Locating hidden servers. In *2006 IEEE Symposium on Security and Privacy (S&P'06)*, pages 15–pp. IEEE, 2006.

[19] İlker Özçelik. *DoS Attack Detection and Mitigation*. PhD thesis, 2015.

[20] İlker Özçelik and Richard Brooks. *Distributed Denial of Service Attacks: Real-world Detection and Mitigation*. CRC Press, 2020.

[21] Paul Pearce, Ben Jones, Frank Li, Roya Ensafi, Nick Feamster, Nick Weaver, and Vern Paxson. Global measurement of {DNS} manipulation. In *26th {USENIX} Security Symposium ({USENIX} Security 17)*, pages 307–323, 2017.

[22] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *Journal of machine learning research*, 12(Oct):2825–2830, 2011.

[23] skivvies. Proxied sites configuration, 2013.

[24] Nicholas Weaver, Robin Sommer, and Vern Paxson. Detecting forged tcp reset packets. In *NDSS*, 2009.

[25] Ryan Whitwam. Supposedly non-existent VPN logs help FBI catch internet stalker, 2017.

[26] Philipp Winter, Tobias Pulls, and Juergen Fuss. Scramblesuit: A polymorphic network protocol to circumvent censorship. In *Proceedings of the 12th ACM workshop on Workshop on privacy in the electronic society*, pages 213–224. ACM, 2013.

[27] Emma Woollacott. If china isn't hijacking internet traffic, there's no reason why not, https://www.forbes.com/sites/emmawoollacott/2018/11/13/if-china-isnt-hijacking-internet-traffic-theres-no-reason-why-not/#453dc9135ed5 (last visited 08/2020), 2018.

[28] Xingsi Zhong, Afshin Ahmadi, Richard Brooks, Ganesh Kumar Venayagamoorthy, Lu Yu, and Yu 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.