

Flexsealing BGP Against Route Leaks: Peerlock Active Measurement and Analysis

Tyler McDaniel, Jared M. Smith, and Max Schuchard

University of Tennessee, Knoxville

{bmcdan16, jms, mschucha}@utk.edu

Abstract—BGP route leaks frequently precipitate serious disruptions to inter-domain routing. These incidents have plagued the Internet for decades while deployment and usability issues cripple efforts to mitigate the problem. Peerlock, presented in 2016, addresses route leaks with a new approach. Peerlock enables filtering agreements between transit providers to protect their own networks without the need for broad cooperation or a trust infrastructure. We outline the Peerlock system and one variant, Peerlock-lite, and conduct live Internet experiments to measure their deployment on the control plane. Our measurements find evidence for significant Peerlock protection between Tier 1 networks in the peering clique, where 48% of potential Peerlock filters are deployed, and reveal that many other networks also deploy filters against Tier 1 leaks. To guide further deployment, we also quantify Peerlock’s impact on route leaks both at currently observed levels and under hypothetical future deployment scenarios via BGP simulation. These experiments reveal present Peerlock deployment restricts Tier 1 leak export to 10% or fewer networks for 40% of simulated leaks. Strategic additional Peerlock-lite deployment at all large ISPs (<1% of all networks), in tandem with Peerlock within the peering clique as deployed, completely mitigates about 80% of simulated Tier 1 route leaks.

I. INTRODUCTION

The Internet consists of many *Autonomous Systems* (ASes) with distinct IP prefixes, routing policies, and inter-AS connections. These networks exchange routes with neighboring ASes over the *control plane* to connect hosts in disparate ASes and create the illusion for users of a single, unified Internet. Unfortunately, there are few security controls on route exchange. ASes behaving adversarially, whether intentionally or by mistake, can export routes that should be kept internally or shared with only a subset of their neighbors. Because the language ASes use to communicate - the *Border Gateway Protocol* or BGP - does not package validation information with routes, remote networks often receive and propagate these *route leaks* throughout the control plane. Leaks frequently steer user traffic on the *data plane* onto unintended paths that lack capacity for the additional traffic. The end result is soaring latency or complete availability loss for destination services. Recent route leaks to prefixes hosting major content/service

providers like Spotify [21], Cloudflare [23] and Google [11] have highlighted the global impact of this problem.

Existing tools designed to curtail leaks, like the many Internet Routing Registries (IRRs), are challenging to deploy or limited in scope. IRRs are databases where ASes can publish their routing policies. Other ASes can then convert IRR-stored policies into filters to validate received routes. IRR-based filtering is limited by its requirement for broad AS participation, however, as the motivations and sophistication of network operators varies greatly between ASes [22]. Other BGP security extensions, like the Resource Public Key Infrastructure (RPKI), only enable filtering for a subset of leaks (e.g. re-origination leaks for RPKI).

The *Peerlock* [44], [18] leak defense system was presented in 2016 to address the need for a deployable solution. Each Peerlock deployment occurs between two neighboring ASes, the *protector* AS and *protected* AS. The protector AS agrees to filter routes that transit the protected AS unless they arrive directly from the protected AS or one of its designated upstreams. The filter prevents the protector AS from propagating or steering its traffic onto any leaked route that transits the protected AS, regardless of origin AS/destination prefix. Peerlock is designed to leverage the rich web of relationships that exist between transit networks in the Internet’s core, and functions without coordination with other ASes on potential leak paths. This makes Peerlock especially viable in the *peering clique* formed by the 19 *Tier 1* ASes that sit atop the inter-domain routing hierarchy. A related technique, Peerlock-lite, enables networks to spot likely leaks without prior out-of-band communication. ASes deploying Peerlock-lite drop routes arriving from customers that contain a Tier 1 AS; it is highly improbable that customers are providing transit for large global networks.

Our first contribution is a measurement of Peerlock/Peerlock-lite deployment on the control plane. In Section IV we design, execute, and evaluate active Internet measurements to search for evidence of filtering consistent with these systems. Our experiments use BGP poisoning, a technique used in prior work for traffic engineering [42] and path discovery [1], to mimic route leaks that transit some target AS. We then listen for which networks propagate - or filter - these "leaks" relative to control advertisements. This information feeds several inference techniques designed to uncover which ASes are Peerlocking for (protecting) the target AS.

Notably, we find substantial Peerlock deployment within the peering clique: about 48% of possible filtering rules (153/342) are already implemented within this set. Further, many non-Tier 1 ASes - including nearly 40% of large ISPs observed during our experiments - perform some Peerlock-lite filtering on Tier 1 AS leaks. Evidence for Peerlock filtering of non-Tier 1 leaks is virtually nonexistent, though three Tier 1 networks (AS 12956, AS 2914, and AS 3320) each filter leaks for more than 20 non-Tier 1 ASes.

After detecting current Peerlock/Peerlock-lite deployments, we ask how well these systems mitigate Tier 1 leaks. Internet-scale BGP simulations in Section V test over 6,000 simulated Tier 1 leaks against observed Peerlock/Peerlock-lite levels to quantify the effect of these systems as deployed. We test the same leaks against six hypothetical extended deployment scenarios to understand where additional filters should be placed to isolate leaks.

We find that Peerlock filtering within the peering clique is helpful, but not sufficient to mitigate Tier 1 route leaks on its own. Consistent with prior work on BGP filtering [8], our experiments show that positioning filters at relatively few ASes - the roughly 600 large ISPs - can play a decisive role in leak prevention. About 80% of simulated leaks were completely mitigated by uniform Peerlock-lite filter deployment at large ISPs, with fewer than 10% of leaks spreading beyond 10% of the topology. These figures are especially encouraging because Peerlock-lite is based on a simple route validity check informed by the valley-free routing model [7] that requires no out-of-band communication.

In this paper, we make the following key contributions:

- We give an overview of the Peerlock and Peerlock-lite filtering systems, and consider their benefits and limitations relative to existing tools in Section III.
- We describe how we adapt existing Internet measurement techniques to probe Peerlock/Peerlock-lite deployment on the control plane and introduce a novel inference method in Section IV-A.
- We actively measure where Peerlock and Peerlock-lite filters are deployed with PEERING [36] and CAIDA’s BGPStream [34] in Section IV-B, with a discussion of results in Section IV-C.
- We simulate thousands of Tier 1 route leaks against several protection scenarios in Section V-B, and present a new path encoding method to understand how these scenarios influence leak propagation and export in Section V-C.

II. BACKGROUND

A. The Border Gateway Protocol

The Internet is a confederation of about 69,000 smaller networks, called *Autonomous Systems* or *ASes*. ASes exchange routing information via the Border Gateway Protocol (BGP) to enable global connectivity. Each AS originates routes to its hosted prefixes; these routes are advertised to neighbors via *BGP updates*. Each update contains a prefix and a collection of other attributes, including an AS PATH that describes the

route’s AS-level hops. ASes compare all received updates via the BGP decision process to select a single best path to every destination prefix. Both path qualities (like AS PATH length) and local network policies (e.g., business relationship with advertising AS) are taken in account when selecting a best path, but policies take precedence in the process. Once an AS selects a best path for a given prefix, it prepends its unique AS number (ASN) to the path and advertises *only* that path to its neighbors.

Paths learned from customer ASes - those purchasing transit - are advertised to all connections. Provider-learned routes, meanwhile, are generally only advertised to an AS’s customers. Peer ASes exchange traffic without compensation, and likewise advertise routes learned from one another only to customer ASes. Limitations on non-customer learned route export prevents customer ASes from transiting traffic between peers/providers at their own expense. This dynamic, known as the Gao-Rexford or *valley-free routing* model [7], guides the exchange of routes on the control plane. No widely-deployed mechanism enforces this model, but the economic incentives it describes shape AS path export behavior.

The *customer cone* [10] is one product of this model. An AS’s customer cone is the set of all ASes reachable from the AS via only provider to customer links. Stated simply, these are the AS’s direct and indirect customers. Customer cone size is one of the few publicly observable features commonly used to judge an AS’s influence on the control plane, e.g. in CAIDA’s AS ranking [50]. Customer cone size is the basis for the UCLA classification presented in [33] widely used in research on this topic [41], [49], [1], [4], [53]. This scheme separates ASes into 1) Tier 1 ASes, who have no providers, form a peering clique, and can transit traffic to any prefix without compensation, 2) large ISPs with more than 50 customer cone ASes, 3) small ISPs with 5-50 customer cone ASes, and 4) stub ASes with fewer than 5 direct or indirect customer networks.

B. Route Leaks

Despite its vital role in binding together Internet networks, BGP is missing key security features like cryptographic hardening of routes exchanged between ASes or trusted certification binding ASes to owned prefixes. This leads to two common classes of major inter-domain routing mishaps, *prefix hijacking* and *route leaks*. Prefix hijacks occur when a network, often unintentionally, originates or advertises a fake but attractive (e.g. shorter or more specific) route to prefixes owned by another AS. Traffic destined for those prefixes is then intercepted by the hijacker. A number of recent studies focus on hijack mitigation [39], [35], [52].

Route leaks are defined in RFC 7908 as the propagation of an advertisement beyond its intended scope [45]. Type 1-4 leaks all cover various valley-free routing violations, i.e. advertising one peer/provider’s routes to another peer/provider. Because remote ASes have little or no information on relationships between non-neighboring networks, they generally cannot distinguish leaks from valid routes, and may propagate

them throughout the topology. Type 5 leaks occur when one provider's routes are announced to another with the AS PATH stripped, effectively re-originating the prefix from the leaker. Finally, a Type 6 leak involves an AS announcing routes used internally to its neighbors. These routes are often more specific than externally announced routes; this makes the leaks more attractive in the BGP decision process and encourages their spread to other remote networks.

Globally disruptive route leaks occur frequently [21], [40], [11], [27], [54], [26]. The 2019 Verizon-Cloudflare leak [23] is a high profile example. A small ISP, AS 33154, leaked specific internal prefixes (Type 6) to Cloudflare and many other destinations to its customer, AS 396531. AS 396531 committed a Type 1 leak by advertising this route to its other provider, AS 701 Verizon. Verizon propagated the leak, which spread widely on the control plane because it was more specific than legitimate available routes (see depiction in Fig. 1). Traffic for Cloudflare, a leading content distribution network (CDN), was funneled through small networks. Many of the thousands of websites and services backed by Cloudflare suffered degraded service until the leak was resolved via out-of-band communication between Cloudflare and AS 33154 [48].

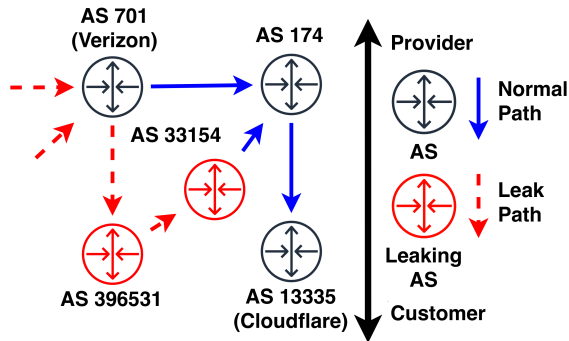


Fig. 1: 2019 Verizon/Cloudflare leak. Other destination services were also affected.

C. Route Leak Prevention

There are a number of tools available to assist network operators in preventing route leaks. The Resource Public Key Infrastructure [25] is a trusted repository for certificates that bind IP prefixes to owning ASes's public keys, called Route Origin Authorizations (ROAs). Remote networks can validate BGP updates against ROAs in the RPKI, a process called Route Origin Validation (ROV). Widespread ROV filtering could prevent Type 5 (and some Type 6) leaks and many prefix hijacking attacks. Unfortunately, ROA/ROV deployment has suffered from circular deployment logic; it is meaningless for origin ASes to invest in publishing ROAs until ROV is widely implemented, but ROV is ineffective without ROAs. This issue has been identified as a major obstacle to ROV deployment [8], [12]. NIST estimates that just 20% of prefixes are covered by a valid ROA [31].

Internet routing registries (IRRs) back another leak prevention system. IRRs are databases where AS operators can store their routing policies. Remote networks can ingest these policies to inform filters that block unintended/invalid advertisements. IRR databases are operated by private firms, regional Internet registries, and other interests [15], and policy entries are often mirrored between them. A complete, up-to-date IRR would eliminate many Type 1-4 route leaks. Like ROV filtering, though, IRR filtering is hampered by deployment headaches. ASes' routing policies are interdependent, so changes to one network's stored policies can render many others obsolete. Operators in smaller, resource-limited networks can avoid periodic updates by configuring permissive routing policies; large transit ASes have complex, dynamic routing policies that require frequent changes to dependent networks' filters [22]. These issues, combined with poor or non-existent authentication, have resulted in inconsistent and out-of-date IRRs. Though leading organizations like RIPE have launched efforts to improve IRR quality [30], operator incentive and dependency issues will continue to limit their usefulness.

Other filtering techniques include max-prefix limit filtering, where a network caps the number of prefixes it will accept from a neighbor. This prevents mass prefix spills like the 2017 Level 3 leak [27], but not more targeted (yet highly disruptive) leaks like the Verizon/Cloudflare incident described earlier. BGPsec [24] is a protocol extension for cryptographic AS path hardening. This would prevent some types of hijacking, but BGPsec has not been commercially implemented and is not designed to prevent route leaks.

Finally, a communities-based "down-only" (DO) leak solution has been proposed [46]. Large BGP communities [43] are signals containing three integers that can be attached to routes. The DO system relies on providers/peers marking a route "down-only" using the first two integers in a large community, with their ASN included as the third integer, before passing the route to customers or peers. If these customers/peers attempt to re-advertise the route to other providers/peers, the attached DO community will clearly signal a route leak. While this system would prevent many leaks if properly implemented, it relies on customers/peers preserving DO information when propagating advertisements. Moreover, some leaks - like the internal route leaks in the Verizon/Cloudflare incident - would not be arrested by this system.

D. BGP Poisoning

BGP poisoning is a technique designed to manipulate the BGP decision process in remote networks. ASes originating a prefix can *poison* an advertisement by including the ASNs of remote networks in the AS PATH. Often, the poisoned ASNs will be inserted between copies of the origin's ASN. This "sandwiching" ensures traffic is routed properly and that the advertisement is valid for ROV filtering purposes (see Fig. 2). BGP prevents cycles from forming in the topology by requiring ASes to drop routes containing their own ASN in the AS PATH; this is known as BGP loop detection. So,

all networks included in the poisoned update's AS PATH - the *poisoned* ASes - will filter it. Poisons can be used for inbound traffic engineering purposes [42], [41], [20] but we employ them in Section IV to mimic route leaks transiting the poisoned AS.

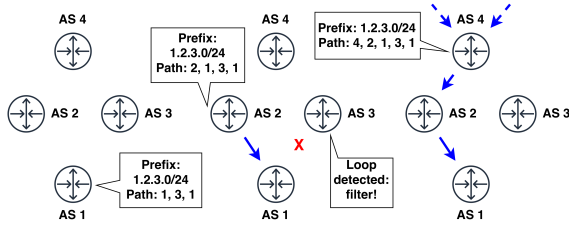


Fig. 2: BGP poisoning. AS 1 originates a route with AS 2 prepended to path (left); AS 2 filters the update (center), but AS 3 propagates (right).

III. THE PEERLOCK SYSTEM

The challenge of leak filtering stems from the topological scope asymmetry between BGP routes and the perspective of individual AS operators who evaluate them. Routes span the topology (global scope); operators only know their own relationships with adjoining ASes (local scope). Filtering systems built on the RPKI [13] and IRRs [14] partially address this asymmetry by applying additional information to the route evaluation process. However, these existing solutions have limitations that have hamstrung their deployment. Most critically, their effectiveness depends on the cooperation of many unincentivized remote ASes as detailed in Section II-C.

A. Peerlock

Peerlock, first detailed by NTT in 2016 [44], [18], is a leak filtering scheme based on out-of-band information exchange between BGP neighbors. Peerlock requires a single AS (the *protected AS*) to designate authorized upstreams to their BGP neighbor (the *protector AS*). This communication distributes AS relationship information between peers to decrease route/filterer scope asymmetry. The protector AS then rejects any BGP update whose AS PATH contains the protected AS unless received 1) directly from the protected AS, or 2) from an authorized upstream, with the protected AS immediately following the authorized upstream in the AS PATH. We say that the protector AS is Peerlocking for the protected AS. See Fig. 3 for a depiction of the system. In this paper, we will often refer to a single instance of Peerlock - that is, one protector/protected pairing - as a *Peerlock rule*.

Here we describe Peerlock's benefits and drawbacks relative to previous leak prevention systems, each of which is described in detail in Section II-C. These comparisons are summarized in Table 1.

RPKI/ROV Comparison: Peerlock provides broader leak type coverage than RPKI/ROV filtering without a trust infrastructure requirement. However, Peerlock only applies to leaks that violate configured topological rules (Types 1-4), so

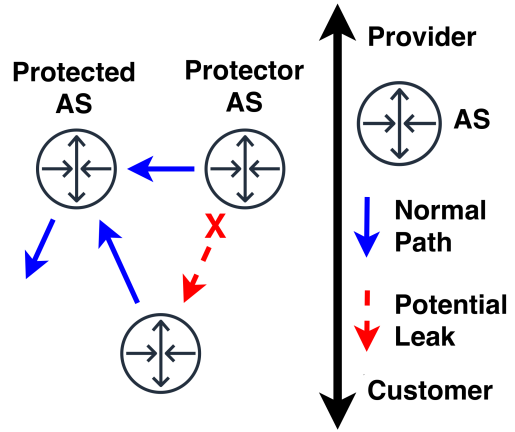


Fig. 3: Simple Peerlock deployment. Protector AS filters updates containing the peer Protected AS from unauthorized propagators.

Type 5 (re-origination) and Type 6 (internal route) leaks fall outside its scope. ROAs tie prefixes to valid originating ASes, so ROV filtering can prevent Type 5 leaks. Additionally, ROAs can be configured with a max prefix length to prevent some internal route leaks and hijacks, although recent work has identified vulnerabilities in this feature [9]. Because Peerlock and RPKI/ROV filtering cover different leak types, Peerlock is *complementary* to ROV filtering rather than a replacement.

IRR Comparison: IRRs are policy object databases capable of storing participating networks' routing intentions with great detail and fine granularity (prefix level). Any AS wishing to enforce these intentions can automatically derive filters from stored objects using software tools, whereas Peerlock rule configuration requires setup between each protector/protected AS pair. Unfortunately, IRRs suffer from incentive misalignment, governance, and rule dependency issues as described in Section II-C. Peerlock rules are self-contained, and changes do not affect other rules. This encapsulation avoids the cascading dependency problem exhibited by IRRs, where one AS's policy changes may render many other AS's entries obsolete.

Most importantly, Peerlock allows the protector AS to stop leaks that transit the protected network regardless of the actions of ASes along potential leak paths; the value of IRR-based filters depends on many remote networks to store accurate policy entries. Peerlock's relatively light cooperation requirement only requires that ASes with an existing relationship communicate information between themselves. This dynamic enables the best resourced, positioned, and incentivized networks (i.e., those serving the most customers) to block route leak propagation regardless of other remote networks' actions.

Max Prefix Comparison: BGP's max prefix feature enables networks to limit the number of prefixes they will accept

System	Coverage	Requirements	Notes
RPKI/ROV	Type 5 internal leak, Type 6 re-origination	RPKI trust infrastructure, local ROA registration & remote ROV checks	Type 5 coverage depends on optional ROA max length
IRR Filtering	Potentially all leak types; depends on stored policy object specificity.	Correct, fresh policy objects & derived filters along potential leak path	Quality issues, misaligned incentives
Max Prefix	All leak types	Filter with meaningful max prefix limit somewhere on potential leak path	Only effective when many prefixes are leaked

Table 1: Common leak filtering systems.

over each neighboring AS connection. Mass route leaks - those involving many prefixes - are filtered once prefix volume over an inter-AS link exceeds the preset limit. Max prefix filtering affords broad protection across leak types, but cannot stop leaks involving few (potentially critical/highly trafficked) prefixes. On the other hand, Peerlock cannot stop leaks that do not violate established topological constraints regardless of volume, but is effective against more selective leaks unprotected by max prefix limits.

Other Considerations: Currently, each Peerlock rule must be manually configured, although at least one method has been proposed to automate Peerlock [16]. Peerlock also lacks a standard to describe how out-of-band information is exchanged between participants. Without a detailed and secure protocol for rule configuration, Peerlock is vulnerable to exploitation; fraudulent rules affect route export, and could be used to engineer traffic flows. Furthermore, operators must define their own ad-hoc protocols for communicating rules that may not guarantee authenticity and/or confidentiality. Virtually all leak solutions discussed here, including IRR, RPKI/ROV, and AS PATH filtering, are recommended by the best practices group Mutually Agreed Norms for Routing Security (MANRS) [5].

B. Peerlock-lite

Peerlock-lite [18] (or Tier 1 filter, "big networks" filter) is a related technique, based on the assumption that transit providers should never receive a route whose AS PATH includes a Tier 1 AS from a customer. This is a valid assumption under the valley-free routing model [7], because such an update implies the customer is providing transit for the Tier 1 AS; otherwise, the customer would not export (leak) the route to another provider. However, Tier 1 ASes have no providers by definition. This logic can be extended heuristically to any other large non-Tier 1 networks that the provider does not expect the customer to export.

This simple logic yields an equally simple filtering rule for transit providers - reject any updates from customers that contain a Tier 1/large transit ASN. See Fig. 4 for a depiction of this filtering technique. Peerlock-lite filters are limited to Tier 1/large transit provider leaks, but they require no out-of-band information to configure. Moreover, Tier 1 ASes' position at the Internet's core results in their frequent presence on AS

PATHs of highly disruptive leaks, e.g. the Verizon/Cloudflare leak [23] and the Enzu/AWS/Spotify leak [21].

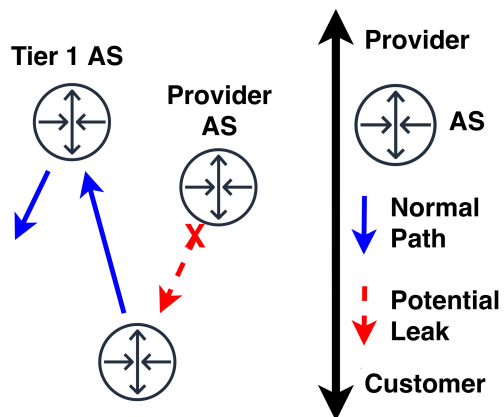


Fig. 4: Example Peerlock-lite deployment. Provider AS filters updates from its customer that include a Tier 1 AS.

IV. MEASURING PEERLOCK DEPLOYMENT

Our initial experiments seek to establish the current state of Peerlock deployment on the control plane. As discussed in the previous section, every Peerlock rule is configured between a pair of networks: the protector AS and the protected AS. Each of the experiments in this section works to identify some or all Peerlock/Peerlock-lite protectors for a targeted AS.

A. Measurement Methodology

Experimental Design: Each set of measurement experiments in this section is designed to discover Peerlock rules for a set of potential protected ASes, called *target ASes*. For each target AS, we advertise a /24 prefix from many points-of-presence (*PoPs*) on the control plane. This is the *control advertisement*. It is a normal /24 origination in every way, except that our university AS - which we know not to be protected by any Peerlock rule - is poisoned (i.e., prepended to the advertisement's AS PATH - see Section II-D). We then listen at varied collection sites, called *collectors*, for BGP updates triggered by our advertisement. The AS PATH for each such update that arrives at collectors lists in encounter order the ASes that received and re-issued the update as described in Section II-A.

Taken together, the gathered AS PATHs form a directed acyclic graph (DAG) that describes the control advertisement’s propagation through the control plane; each AS appearing on at least one AS PATH forms a node in the DAG, and AS ordering within paths allows us to form directed edges between nodes. BGP loop detection prevents cycles as noted in Section II-D. We call this graph the *control DAG*. Note that all of the ASes appearing in the control DAG propagated (and thus did not filter) control updates that include a poisoned AS.

We then wait 30 minutes for update propagation before issuing an explicit withdrawal for the /24 prefix. This timing is built conservatively from empirical measurements of propagation times through the control plane (see update propagation experiments in appendix). After another waiting period to ensure the withdrawal has completely propagated, we issue a *leak advertisement* for the same /24 prefix. This advertisement matches the control advertisement in every way, except that the target AS is poisoned. This leak advertisement structure is designed to mimic a leak for the purposes of Peerlock while avoiding other common filtering systems. The target AS’s presence on update paths triggers filtering for any Peerlock protector ASes.

Finally, we gather all BGP updates for the leak advertisement from our collectors. The ASes that appear on AS PATHs in any of these updates are added to a set called the *leak set*. Since they propagated poisoned updates, we know these ASes did not filter the "leak". With the control DAG and leak set together, we can reason about which ASes are Peerlocking for the target AS using two techniques: 1) clique inference and 2) DAG inference.

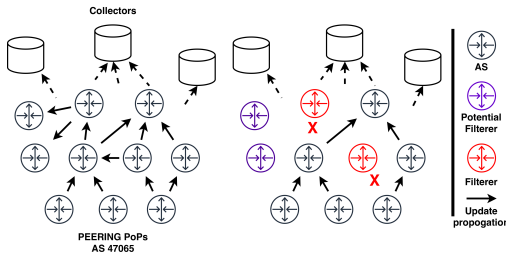


Fig. 5: Measurement experiment depiction. Inferences are made about Peerlock deployment based on differences between normal updates (left) vs. poisoned updates (right) arriving at collectors.

For detecting Tier 1 protector ASes, we use *clique inference*. This simple technique relies on the fact that Tier 1 ASes form a peering clique by definition. According to the valley-free routing model [7], ASes share all updates received from customers with their peers; this maximizes the traffic the AS transits for its customers (and thus the AS operator’s compensation). Further, ASes should not share a peer’s updates with another peer, as this is a Type 2 route leak [45]. So, in general, if a Tier 1 AS is observed propagating an update, *all* Tier 1s should receive the update via their peering relationships. Because we observe at least one Tier 1 propagating control and leak updates across all experiments, we define a simple

rule for inferring Tier 1 protector ASes: any Tier 1 AS that appears in the control DAG but not the leak set is Peerlocking for the target AS.

Inferring other protector ASes requires a more general technique. Outside the structural guarantees provided by the Tier 1 clique, there is significantly more uncertainty about which networks are filtering leak updates. Specifically, it is difficult to distinguish an AS filtering updates from an AS not receiving updates at all due to filtering by other upstream/downstream networks. This challenge leads us to make three separate inferences for these ASes for each leak target.

First and simplest is the *max inference* set, defined as all control DAG ASes minus the leak set. This set includes all ASes who *may have* filtered leak updates, but also ASes who did not receive the leak update because it was filtered by an intermediate AS. Secondly, we build a *min inference* set. This set is built by deleting all leak set ASes from the control DAG, and collecting the root of every weakly connected component that remains. This isolates the ASes that filtered leak updates from ASes in their "shadow" who did not receive the updates. The min inference set contains those ASes who likely filtered leak updates based on routes we observed. Note that the min/max inference techniques closely align with those employed in the long path filtering experiments in Smith et al.’s study on BGP poisoning as a re-routing primitive [41].

Our last inference set is the *likely inference*. Because ASes only export their best path to our /24 prefix, we cannot observe every edge that should exist in the control DAG (i.e., every potential propagation path for updates). So, this set’s is built like the min inference set, except that we augment the control DAG with edges from CAIDA’s provider-peer observed customer cone inference [51]. That is, we add edges to the control DAG where CAIDA’s data indicates there are links between ASes that we did not observe due to policy decisions. This forms a superset of the min inference set and a subset of the max inference set that contains the most likely filterers. This is a novel technique not used to our knowledge in any prior work on this topic.

These three inference sets are formed for each target from differences in control and leak update propagation. In addition to these sets, we also build a *min/max/likely poison filtering set* by following the same steps listed above, but with a *unpoisoned advertisement’s* updates compared against the control advertisement’s updates. These sets are built to explore the prevalence of general poison filtering as in Smith et al. [41].

Framework Details: The control-plane measurement framework for these experiments consists of 1) 13 PoPs to issue BGP advertisements and 2) 54 BGP collectors to listen for propagation. We employ the PEERING testbed [36] for the first requirement. PEERING allows us to advertise three assigned /24 prefixes from edge routers at thirteen PoPs worldwide. For collecting BGP updates, we used CAIDA’s BGPStream [51] tool. This tool draws updates from 54 globally distributed collectors, including 30 RouteViews [32]

and 24 RIPE RIS [29] collectors. While most of these collectors are positioned in North America and Europe, every populated continent is represented by at least one collector.

Measurement Limitations: While our framework allows us to effectively probe the control plane for evidence of Peerlock and related techniques, a number of limitations prevent complete certainty regarding Peerlock filter placement. The most important of these obstacles are imperfect collector coverage, topological instability, and the presence of other filtering systems. Here we discuss each of these factors in turn.

BGP policies prohibit us from viewing the entirety of the topology with our framework; there are few collectors in stub networks, and stub/remote ASes do not export received updates back "up" through provider networks. This means our *observation window* - the ASes on update paths at collectors - is biased toward transit networks in the Internet's core as in [33]. Fortunately, this is the most important/influential region to monitor, as these network's policies have the widest impact on the control plane. Altogether, we observed 610 ASes during our experiments, including 181/605 large ISPs and all 19 Tier 1 networks. Most observed ASes (332) were present in the observation window during all experiments conducted from August 2019-May 2020. Note that while we can only infer *protector* ASes from our observation window, we can poison any AS. So, our window does not limit our inference regarding which ASes are *protected*.

To account for instability in our observation window, we limit our filtering inferences to those ASes observed in control updates both before and after the leak advertisement (i.e., for the current and next target AS experiment). Additionally, we repeat experiments - issue control/leak advertisements for the same target ASes - over several months. These observations are combined to reduce the "noise" of topological dynamism from our inferences. Specifically, we remove ASes from a target's filtering inference sets if we later observe them propagating a leak update for that target; in this case, the earlier inference was likely caused by the AS's intermittent presence in the observation window during the experiment.

Most importantly, we acknowledge that we cannot be certain Peerlock/Peerlock-lite exactly as described by NTT [44], [18] is responsible for all observed filtering, but our experiments are designed to avoid common leak filtering systems. First, since the leak and control advertisements in our experiments share an origin AS/prefix, their updates present identically for ROV filtering purposes. Additionally, since we observe all ASes in the control DAG propagating control updates, we infer those ASes will not apply common IRR or max-prefix limit filtering to the same /24 in leak updates. Finally, while prior work indicates that short poisoned paths are frequently present on the control plane [49] and rarely filtered [41], the poisoning in the control advertisement ensures that we do not conflate poison filtering and Peerlocking.

Despite our efforts to avoid common filtering practices, local policies grant network operators extensive discretion in

how routes are vetted and exported. This flexibility means we cannot be certain that experimental updates are not sometimes blocked by AS specific, ad-hoc AS PATH filtering techniques. We know of no way to distinguish such functionally similar filters from Peerlock.

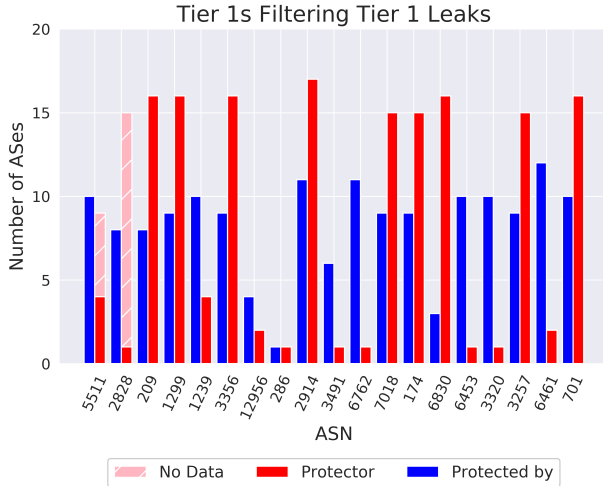
Ethics: We issued only well-formed BGP advertisements using the PEERING software client and adhered to all rules published by PEERING. We advertised only our assigned /24 prefixes, which are reserved for experimental use, and thus did not disturb Internet control or data plane operation for any non-experimental IP addresses. Our experiments did require poisoned advertisements, but this is a common practice used both in research [2], [41] and in traffic engineering [49]. One network operator observed and inquired about our experiments to PEERING, but did not report any resultant adverse effects. No data-plane traffic was sent during the conduct of our experiments.

B. Evaluation

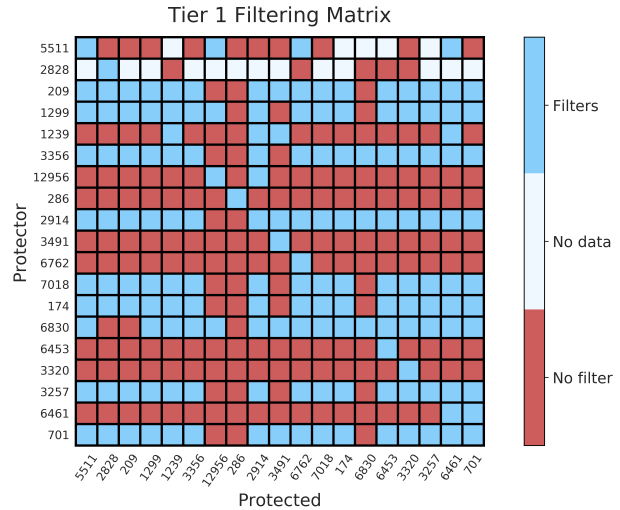
Target Set 1, Tier 1s: The 19 Tier 1 ASes form our first target AS set, i.e. the potential protected ASes for which we are inferring Peerlock rules. The Tier 1 peering clique includes the most influential networks by one of the few observable metrics, customer cone size [51], and often creates [27] or distributes [23], [11], [40] leaks that disrupt global Internet services. Paradoxically, deploying filters for leaks that include Tier 1 ASes is also relatively simple for non-Tier 1 networks via the Peerlock-lite system described above. We iteratively issued unpoisoned, control, and leak advertisements that covered this set every two months from August 2019 to May 2020. This repetition allows us to capture filtering rules for ASes with inconsistent presence in our observation window, and to explore how deployments change over time.

We first present results for protection within the Tier 1 clique in Fig. 6. Note that because of BGP loop detection, every AS filters leak updates that include their own ASN regardless of Peerlock deployment. The peering clique is fortunately the most stable feature in our observation window, enabling us to measure the presence/absence of nearly every potential Peerlock rule within the clique. We have marked the exceptions for which we were unable to measure filtering rules in pink in Fig. 6. We see that Peerlock deployment is significant but unevenly distributed within the clique. Some ASes - e.g. AS 2914 NTT, AS 701 Verizon - filter leak updates for virtually the entire clique. For five others - e.g. AS 3491 PCCW Global, AS 6762 Telecom Italia - we found no evidence of Tier 1 Peerlock filtering at all.

Our measurement results for Peerlock/Peerlock-lite protection of Tier 1s by all observed ASes are depicted in Fig. 7. Fig 7a shows both our inferences about which networks filter poisoned updates in general (blue lines) and which filter Tier 1 leaks (red lines). These are displayed as a cumulative distribution function (CDF) over Tier 1 targets; likely inferred filtering levels range from about 3% (AS 6830) to 15% (AS 701) of observed ASes. Note that per



(a) Number of protector/protected rules by ASN. Protector numbers include ASes protecting their own ASN via loop detection.



(b) Depiction of Tier 1 protection rules.

Fig. 6: Tier 1s filtering Tier 1 leaks, 2019/2020 measurements.

the experimental design described above, we cannot make Peerlock protection inferences for ASes filtering all poisoned updates; however, this is a small set without Tier 1/large ISP members (max inference size = 9 ASes). Fig. 7 shows the number of ASes in each UCLA class (see Section II) protecting at least one Tier 1 target.

Target Set 2, Tier 1 Peers: Our second target set includes the non-Tier 1 peers of Tier 1 ASes (about 600 ASes) as inferred by CAIDA [51]. These experiments explore whether Tier 1 ASes are extending Peerlock protection to their non-Tier 1 peers. Additionally, despite covering about 1% of all ASes, this set includes a third of all large ISPs. The presence of these large ISPs in the target set allows us to investigate whether non-Tier 1 ASes apply Peerlock-lite filters to large transit networks outside the peering clique. These experiments were conducted from Oct 2019 to May 2020, with every included network targeted at least twice.

The overall results are presented in Fig. 8a. Clearly, filtering for these leaks is less prevalent within our observation window. 80% of Tier 1 peer leaks were filtered by fewer than 2% of observed ASes, but a few exceptional targets did trigger significant filtering behavior. Our poison filtering inference for these targets is, as expected, nearly identical to that derived from the Tier 1 leak experiments. Fig. 8b displays filtering levels for each Tier 1 ASes by peering status with the target. All Tier 1s protect 10 or fewer peer networks from this set. More variance exists in non-peer filtering behavior, as we will explore in the following discussion.

C. Discussion

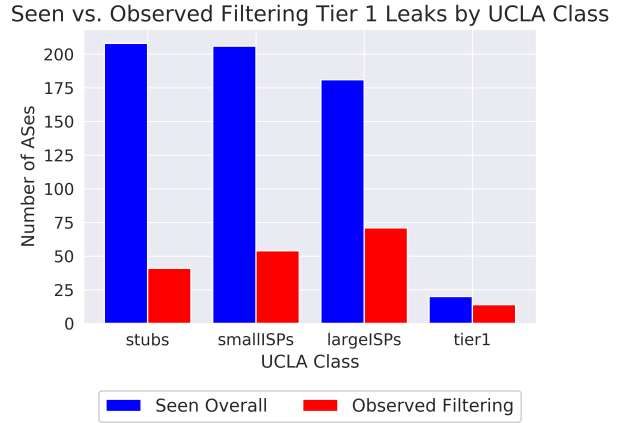
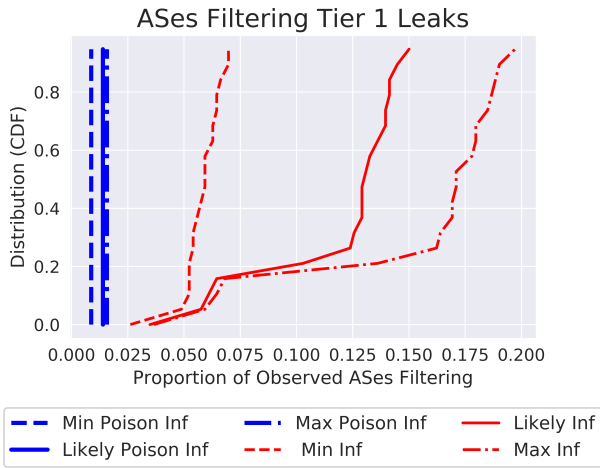
Consistent with Smith et al. [41], we find no evidence for widespread filtering of otherwise unremarkable poisoned paths. Their study also found that poisoning high degree ASes in an update is associated with reduced propagation.

Specifically, sub-20% update propagation rates were observed for some Tier 1 ASes, including AS 174 (Cogent/Tier 1) and AS 3356 (Level 3/Tier 1). Birge-Lee et al. [2] likewise found that using AS poisoning rather than communities as a path export control primitive significantly reduced update spread, especially when large transit providers were poisoned. Defensive AS-path filtering (e.g., Peerlock/Peerlock-lite) is identified as a likely culprit for this effect. Our work systematically examines how and where these filters are deployed on the control plane (within the limits of our observation window).

Tier 1 Leak Filtering: The greatest protection within our observation window is clearly afforded to Tier 1 ASes. Our initial experiments in August 2019 discovered evidence for 133/342 ($19^2 - 19$) possible Tier 1-Tier 1 filtering rules (about 39%). Each measurement that followed uncovered at least two new filtering rules, and by our final experiment in May 2020, 153 rules had been observed, a nearly 15% uptick in Peerlock deployment. We had previously observed a negative filtering result for every additional rule, indicating this increase results from genuinely new Peerlock deployments rather than instability in the observation window.

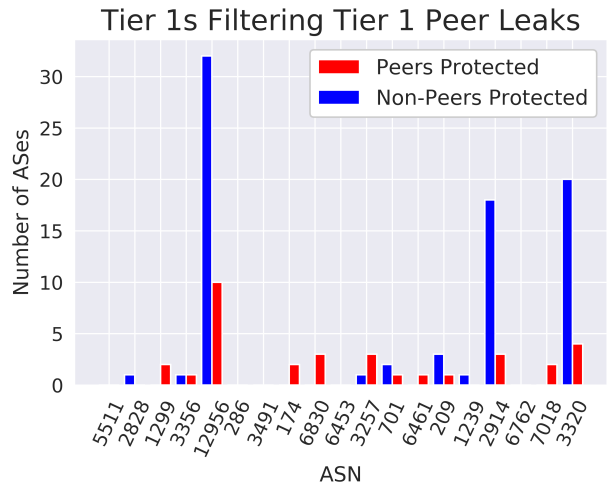
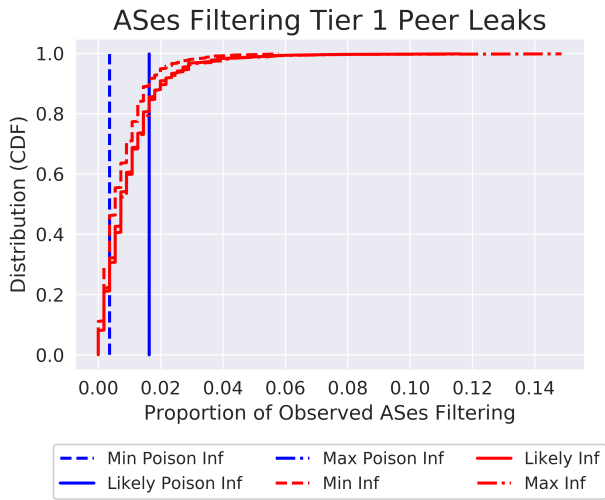
Non-Tier 1 ASes also filter Tier 1 leaks, though this behavior is far from uniform. Overall, Tier 1 leak filtering ranged from 3% to 15% of observed ASes across Tier 1 AS targets. Most of this is likely due to Peerlock-lite filtering, as it is simpler to deploy. Moreover, fewer than 10% of the more than 1,000 observed Tier 1 filtering rules exist between peers, and only about 20% (236 rules) involved a Tier 1’s indirect customers filtering leak updates. This suggests that ASes are installing Peerlock-lite filters for all Tier 1s rather than simply protecting their upstream providers.

Mutually Agreed Norms for Routing Security (MANRS) [6] is an initiative whose ISP members agree to best routing prac-



(a) Blue lines show poison filtering; red lines depict Tier 1 leak filtering. (b) Blue bars show no. ASes in observation window; red bars show no. ASes filtering at least 1 Tier 1 leak.

Fig. 7: Overall filtering of Tier 1 leaks, 2019/2020 measurements.



(a) Overall filtering levels for Tier 1 peer leaks. Max and likely poison inferences match for this set. (b) Tier 1 filtering of Tier 1 peer leaks (peers within clique excluded).

Fig. 8: Tier 1 peer leaks, 2019/2020 measurements.

tices (like AS path filtering) to secure inter-domain routing. While Peerlock and Peerlock-lite are not specifically included in MANRS expected filtering actions, they are both suggested in the implementation guide [5]. Fig.9 displays as a CDF the proportion of MANRS and non-MANRS ASes filtering Tier 1 leaks. 73 of 502 MANRS ASes fall within our observation window; the proportion of observed MANRS ASes that filtered Tier 1 leaks ranged from 2-18% depending on Tier 1 target. Non-MANRS filtering over the same target set ranged from 2 to 12%.

As shown in Fig 7b, the proportion of ASes with Tier 1 leak filters rises with UCLA class [33]. Intuitively, networks with larger customer cones have the resources for sophisticated configurations and the imperative to prevent

issues for downstream customers, and have previously been associated with differing responses to BGP events [41], [3]. This dynamic hampers systems requiring wide participation like ROV [8] and IRR filtering [22], but does not limit Peerlock or Peerlock-lite deployment.

Tier 1 Peer Leak Filtering: Our non-Tier 1 leak experiments met with relatively sporadic filtering. For more than 80% of targets in this set, nearly every observation window AS ($\geq 98\%$) propagated leaks. As described in Section III, Peerlock-lite filters for non-Tier 1 ASes require more careful deployment. The outliers in this target set (see the long tail in Fig. 8a) are invariably near-Tier 1 networks like AS 1273 Vodafone, AS 6939 Hurricane Electric, and AS 7843 Charter

that are safe for most ASes to include in a Peerlock-lite filter.

Tier 1 filtering of this leak set was likewise reduced compared to Tier 1 leaks. In general, Tier 1 networks deploy fewer than 5 Peerlock filters for non-clique peers. Nearly all of these cover near-Tier 1s like AS 7922 Comcast and AS 1273 Vodafone, or ASes administered by Tier 1s e.g. AS 702/703 Verizon and AS 3549 Level 3. Notably, three networks extend protection to more than 15 non-peers (per CAIDA’s inference). AS 2914 NTT’s non-peer filtering rules all cover various Comcast ASNs, while AS 12956 Telefonica’s rules appear to be regionally-based: zero rules are applied to customer cone ASes, but 23/31 apply to other European ISPs of varying size. 13/20 of AS 3320 Deutsche Telecom’s non-peer filtering rules, on the other hand, cover ASes within its customer cone.

In summary, Peerlock is widely deployed and expanding within the peering clique. Deployment outside the peering clique is relatively limited, however. Up to 20% of non-clique networks also deploy Peerlock-lite (or a similar mechanism) to filter leaks containing Tier 1 or near-Tier 1 ASes. These deployments are proportionally more common in ISPs and rarely seen in stub ASes within our observation window. Fortunately, the effectiveness of Peerlock/Peerlock-lite deployments is less sensitive to scattershot deployment than other filtering solutions. Prior work [8] and our simulations in the following Section V suggest that filtering by large ISPs can have an outside impact on global leak propagation.

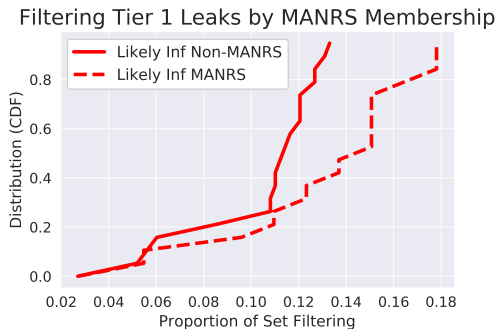


Fig. 9: Tier 1 leak filtering for MANRS/non-MANRS ASes.

V. EXPLORING PEERLOCK’S PRACTICAL IMPACT

The substantial but limited Peerlock/Peerlock-lite filtering measured in the previous section leads us to investigate these systems’ protective benefit in partial deployment. We have interest both in how well these systems protect the control plane from Tier 1 leaks as deployed, and in the relative effect of realistic additional deployment (e.g. adding filters at large transit networks). To answer these questions, we quantify Peerlock’s practical impact with Internet-scale leak simulations against several filter deployment schemes.

A. Simulation Methodology

These experiments are conducted via extensions to a BGP simulator, an approach consistent with prior work on this

topic [38], [37], [42], [49]. We construct a simulated AS-level topology from CAIDA’s inferred relationship dataset (Jan. 2020 data) [51]. ASes within the topology evaluate and export routes using the BGP decision process; longest-prefix matching, LOCAL PREF, and AS PATH guide path selection, while route export is governed by local policy to enforce valley-free routing. This ensures the simulator models the central dynamic of control plane propagation - the Gao-Rexford model [7], and allows for the closest approximation of control plane behavior we can devise without ASes’ full (private) routing policies.

Each simulation is driven by a *protection scenario* that maps protector ASes to those they are protecting. As with Peerlock in practice, these protectors drop all received routes that transit a protected AS unless they arrive directly from that AS. Some scenarios also include Peerlock-lite deployments; for these experiments, some set of ASes filter all customer-exported routes that transit Tier 1 ASes. Once we establish the protection scenario, we iterate over all Tier 1 to Tier 1 links (with 19 Tier 1 ASes, this is $n = 19, n^2 - n = 342$ links). These links describe a unidirectional connection from one Tier 1, called the *link start*, to another Tier 1, called the *link end*.

For every link in this set, we sample 20 ASes from the link start’s customer cone to serve as *leakers*. Each leaker will, in turn, randomly select a *destination AS* in the link end’s customer cone, and advertise a route to the destination over the link to all of its peers/providers (see Fig. 10). This models a Type 1 route leak of a path over the peering clique [45]. After the leak, we allow the topology to converge and measure how many ASes 1) received leak updates and 2) installed the leak path. Additionally, we capture all the AS PATH of all leak updates for analysis. With 20 leaker/destination pairings per link and 342 Tier 1 links, we simulate 6,840 leaks in total.

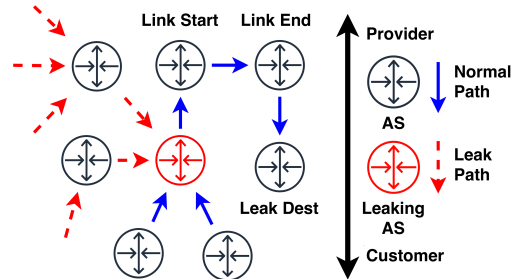
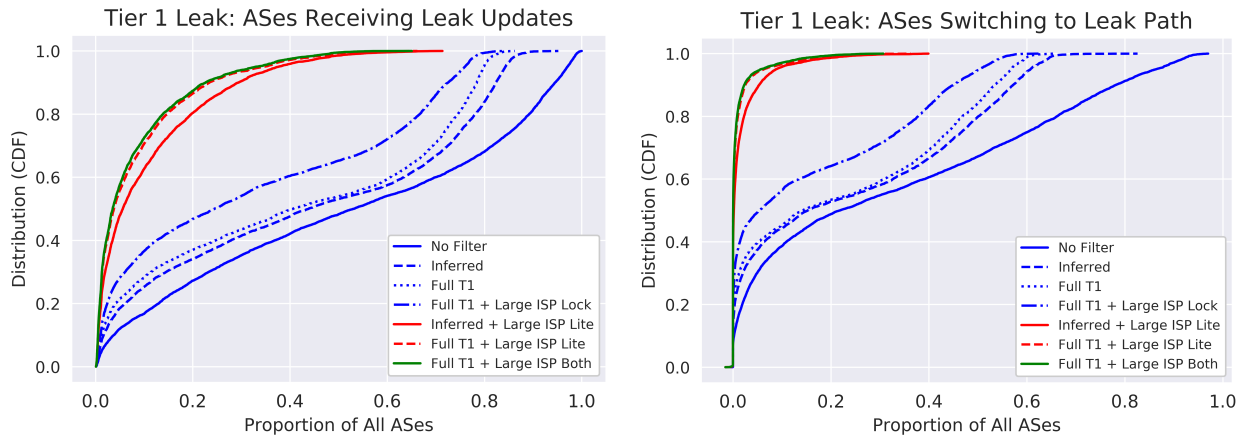


Fig. 10: Example simulated leak. Dashed red lines indicate route leak to other providers/peers.

Our simulations focus on leaks with Tier 1 leaks for two reasons. First, we do not find substantial real-world Peerlock/Peerlock-lite protection of non-Tier 1 ASes as outlined in Section IV. Second, many consequential leaks are propagated globally over the Tier 1 backbone, e.g. [27], [23], [40], [11]. Some of our protection schemes will investigate whether leaks can propagate throughout the Internet without Tier 1 distribution.



(a) Impact of various deployment scenarios on leak update propagation. (b) Note increased Peerlock-lite performance for path switching vs. leak update propagation.

Fig. 11: Peerlock/Peerlock-lite simulation results.

B. Evaluation

We evaluate seven different protection schemes for Tier 1 leaks.

- **No filters.**
- **Inferred:** Tier 1 Peerlock levels observed during Internet measurements.
- **Full T1:** All Tier 1s Peerlock for all other Tier 1s.
- **Full T1 + large ISP lock:** Same as full T1, but all large ISPs (376 ASes in CAIDA Jan 2020 dataset [51]) Peerlock their Tier 1 peers.
- **Full T1 + large ISP lite:** Same as full T1, but all large ISPs deploy Peerlock-lite to protect clique ASes.
- **Full T1 + large ISP both:** Same as full T1, but all large ISPs deploy Peerlock-lite filters and Peerlock for their Tier 1 peers.
- **Inferred + large ISP lite:** Same as inferred, but all large ISPs deploy Peerlock-lite.

While it is simpler to filter customer-learned routes with Peerlock-lite than to deploy Tier 1 Peerlock filters for large ISPs, we include both Peerlock and Peerlock-lite filtering by these ASes to study how leaks are propagated within the topology. The results of these experiments are presented in Fig. 11, which displays both the proportion of ASes in the topology receiving leak updates (Fig. 11a), and the proportion selecting/exporting the leak path (Fig. 11b).

A critical feature revealed by Fig. 11 is the insufficiency of Tier 1 protection alone (blue lines). Full Tier 1 Peerlocking prevents all distribution of studied leaks over the peering clique, but leak updates still spread to the majority of the topology for most experiments. Adding large ISP Peerlock protection has a relatively significant impact on both propagation and installation.

Peerlock-lite deployment by these ASes (red lines) benefits from more filterers with wider protection per filterer. Naturally, these scenarios are much more effective at preventing prop-

agation. For most leak cases, less than 10% of the topology receives leak updates. This highlights the leverage large ISPs have within the topology; filtering at these ASes (<1% of all networks) generates an extensive shielding effect. The distinct "shoulder" on the Peerlock-lite curves in Fig. 11b suggests the impact on ASes *using* the leak is even more pronounced. There is virtually no impact on target link usage for 75% of simulated leaks when Peerlock-lite is deployed by all large ISPs. Interestingly, the combination of Peerlock and Peerlock-lite filtering by large ISPs (green line) adds little value over Peerlock-lite alone.

C. Discussion

Path Encoding: To analyze how each of these scenarios shapes leak propagation (and route selection/export), we collect the AS PATH of all leaks exported during the above experiments. We use a novel *path encoding* whereby each AS on leak AS PATHs is converted to a 2-tuple with the form (relationship to next AS, UCLA class [33]). Only the AS PATH segment from the first customer to provider link to the leaker ASN - the *leak segment* - is encoded. This trimming discards the "down" segment prepended as leaks propagate within customer cones, as well as the the segment connecting leaker and destination that is invariant across leaks. We include AS relationship in the encoding because of its importance in path export behavior as described in Section II-A; UCLA class informs us regarding where leaks travel through the routing hierarchy. Taken together, these factors help us understand broadly the topological dynamics at play in leak propagation, and to capture the dominant leak propagation vectors under each protection scenario.

Relationship is encoded as "C" (customer), "R" (peer), or "P" (provider). UCLA classes are indicated by "T" (Tier 1), "L" (large ISP), "S" (small ISP), and "U" (stub). Example: [LR, TP] encodes a leak path exported to a Tier 1 provider by the leaker, who then passes the leak to a large ISP peer. The

Common Leak Segment Encodings		
Scenario/Encoding	No. exporting ASes	% of exporting ASes
No filters	141,797,992	100%
[LR, LP]	14,892,311	11%
[TP]	10,254,707	7%
[LR, LP, LP]	8,683,968	6%
Inferred	108,030,704	100%
[LP, LP]	14,325,960	13%
[LR, LP, LP]	8,675,841	8%
[SR, LP]	5,169,427	5%
Full T1	101,024,444	100%
[LR, LP]	14,246,024	14%
[LR, LP, LP]	8,978,175	9%
[SR, LP, LP]	5,163,786	5%
Full T1 + large ISP lock	69,638,282	100%
[LR, LP]	9,473,820	14%
[LR, LP, LP]	5,899,779	8%
[SR, LP]	3,310,842	5%
Inferred + large ISP lite	8,005,724	100%
[LR]	2,537,276	32%
[LR, TP]	1,281,620	16%
[LR, SP]	653,167	8%
Full T1 + large ISP lite	5,215,232	100%
[LR]	2,386,597	46%
[LR, SP]	679,076	13%
[SR]	412,399	8%
Full T1 + large ISP both	4,649,828	100%
[LR]	2,023,579	44%
[LR, SP]	584,124	13%
[SR]	407,661	9%

Table 2: Most common encodings with number and percentage of ASes exporting leaks.

progress of the leak through the large ISP’s customer cone would continue to the left of "LR", and the path from leaker to destination would continue to the right of "TP", but these segments are omitted as explained above.

We will use two tables in analyzing our results. Table 1 depicts the three most common leak encodings for each scenario; these account for at least a quarter of leak paths regardless of filter placement. We also list the sum and percentage of ASes exporting leaks accounted for by each encoding. Table 2 gives summary statistics for leak segments, including their average length and the percentage of leak segments transiting each UCLA class. Because we do not encode customer cone propagation in leak segments, stubs are transited in <10% of paths across all protection scenarios, are omitted from the table.

First, we observe that even under the "no filters" scenario, leaks re-transiting the Tier 1 clique are not the most common path encoding in Table 1. Table 2 shows they are present in <35% of leak segments under all scenarios. This result is an artifact of the BGP decision process; paths learned from customers are preferred over those exported by peers, and peer routes are favored over provider-learned ones. So, with all other selection criteria equal, routes exported from providers "above" an AS in the topology - e.g. the peering clique - will generally only be installed and exported if the AS has not received an update from peers/customers "below". Since Tier 1 providers cap the routing hierarchy, we expect ASes will prefer non-Tier 1 routes when provided alternatives by their connectivity. This dynamic explains why the additional

protection afforded by complete Peerlock within the peering clique vs. current levels is muted in Fig. 11b.

This effect also brings large ISPs to the fore in our simulations. As noted in [33], these networks are densely connected with peering links. Their connectivity allows them to bypass the Tier 1 clique for many routes - and makes them the primary channel for leak propagation. The most common encoding for every scenario in Table 1 includes a large ISP, and 18/21 of the top encodings transit at least one. More than 70% of leak segments transit these ASes for all protection scenarios (see Table 2). In fact, in the scenarios without Peerlock-lite (top four listed), leak segments on average transit - and could be filtered by - multiple large ISPs. These statistics motivate the scenarios that place Peerlock-lite filtering at these ASes (bottom three in tables).

Interestingly, Peerlock-lite diminishes leak usage and propagation unequally as shown in Fig. 11. Fig. 11a shows about 20% of leak segments propagate to 20% or more of the topology with large ISP Peerlock-lite deployment, but Fig. 11b shows that fewer than 5% are installed/exported by at least 20% of ASes. Table 1 hints at why this is the case - a third or more of leak segments in Peerlock-lite scenarios are exported to large ISP peers, who propagate them directly into their customer cones (indicated by [LR]). Large ISPs with any customer-learned or preferential (e.g. shorter) peer-learned paths to the leak destination will prefer their existing route, so the [LR] only includes a subset of the leaker’s peers. Large ISP peers advertising the leak to customers could reach many ASes, but as a provider-learned route, the leak will be disadvantaged in the BGP decision process.

We see in Table 1 and Fig. 11 that small ISPs do not have the connectivity to propagate leaks globally when the large ISP provider channel is blocked by Peerlock-lite. Under all scenarios, most leak segments do not transit a small ISP (though they may be transited during propagation into customer cones). This feature suggests a less prominent role in route exchange for these networks relative to large ISPs.

To summarize, we find large ISPs are the most critical players in halting the spread and installation of Tier 1 leaks. These networks are interconnected enough to globally disseminate route leaks *without the peering clique* in many cases. Moreover, adding simple Peerlock-lite filters at these ASes to the currently deployed Peerlock filters in the peering clique causes a 94% reduction in total leak export across 6,840 leak simulations. Table 1 suggests that peer connections among ISPs are the largest remaining vulnerability for Tier 1 leaks given uniform large ISP Peerlock-lite deployment. These channels are out of reach for Peerlock-lite as described, but could be mitigated by 1) additional peering relationships/Peerlock rules to protect important leak targets and/or 2) complementary leak prevention systems like IRR filtering.

VI. RELATED WORK

Smith et al.’s 2020 study on the efficacy of poison filtering for inbound re-routing [41] similarly employed the PEERING framework to probe the behavior of remote networks. That

Transited AS Statistics												
Scenario	Segment Length			Tier 1s			Large ISPs			Small ISPs		
	average	std. dev	% paths	average	std. dev	% paths	average	std. dev	% paths	average	std. dev	
No filters	4.4	1.8	27%	0.2	0.4	89%	2.3	1.6	40%	0.7	0.8	
Inferred	4.6	1.8	7%	0.5	0.2	98%	2.5	1.5	45%	0.8	0.9	
Full T1	4.7	1.8	0%	0.0	0.0	98%	2.6	1.5	46%	0.8	0.9	
Full T1 + large ISP lock	4.8	2.1	0%	0.0	0.0	98%	2.7	1.6	50%	0.9	1.0	
Inferred + large ISP lite	2.9	1.0	35%	0.3	0.5	72%	0.6	0.5	42%	0.8	1.0	
Full T1 + large ISP lite	2.7	1.1	0%	0.0	0.0	74%	0.6	0.5	47%	1.0	1.1	
Full T1 + large ISP both	2.7	1.1	0%	0.0	0.0	71%	0.6	0.5	49%	1.0	1.1	

Table 3: Analyzing leak segments by UCLA classes transited.

work encountered some evidence for poison filtering, and noted that filtering rates increase with poisoned AS degree, but did not seek to describe the underlying filtering mechanism or measure which ASes filter poisons. Similarly, Birge-Lee et al. [2] and McDaniel et al. [28] trialed poisons as primitives for novel BGP attacks. Both studies encountered filtering when attempting to poison large transit networks, but did not examine filtering position or prevalence.

Hlavacek et al. [12] introduced the DISCO system for preventing BGP hijacking. While not designed to prevent route leaks, the approach taken by DISCO is ideologically similar to Peerlock - DISCO emphasizes deployability/usability at the expense of some security guarantees relative to RPKI/ROV filtering. This line of thinking is informed by a long history of glacial deployment rates for security features that harden BGP including BGPsec [47] and the RPKI [8], [31].

Previous work that relies on BGP poisoning often assumes 1) unpoisoned ASes will forward poisoned updates and 2) poisoned ASes will drop such updates (see Section II-D). Katz-Bassett et al.’s LIFEGUARD fault detection and remediation system [19], [20], for instance, employs poisoning to steer ASes around link failures. Smith and Schuchard’s Nyx defense [42] depends on rerouting with poisons for DDOS/Link Flooding Attack mitigation. Anwar et al.’s path discovery technique [1] is also driven by BGP poisoning. While we discovered little evidence for general poison filtering, the prevalence of Tier 1/large transit network filtering could present an obstacle to these systems. Specifically, the assumption that unpoisoned networks will propagate poisons does not hold in all cases.

VII. CONCLUSION

This work probes the current deployment of Peerlock/Peerlock-lite on the control plane with active Internet measurements in Section IV. We find substantial evidence for deployment of these leak defense systems, especially in large transit networks, and measure a rise in Peerlock deployment within the peering clique during our experiments. While the range of protected networks is still narrow within our observation window, with most filterers protecting only Tier 1 ASes, many of the most disruptive recent route leaks contain these networks. Defensive systems [42], [20], measurement techniques [1], and attacks [2] that may poison Peerlock-protected networks could inadvertently trigger these

filters, and must not assume poisons will be propagated by unpoisoned networks. BGP simulators should likewise account for the presence of Peerlock to faithfully reproduce control plane behavior.

We also examine how the position and prevalence of filtering impacts leak propagation in the AS-level topology in Section V. Notably, we find that large ISPs filtering plays a major role in global leak dissemination, signaling that Tier 1 clique deployment of Peerlock alone is not sufficient to isolate leaks. Strategic placement of filters at these large transit providers, which account for fewer than 1% of all ASes, completely mitigates 80% of simulated Tier 1 route leaks.

The MANRS filtering guide encourages AS PATH filtering by member ISPs, particularly for screening customer advertisements, and gives Peerlock/Peerlock-lite as examples. But these systems are not explicitly required, unlike IRR filtering (see [5] Section 4.1.1.1). Given the many indirect/direct customers these networks serve, ISPs are best equipped and best incentivized to deploy effective filters. Moreover, neither Peerlock nor Peerlock-lite is technically complex or burdensome to configure. Therefore, we argue for broad application of these common-sense leak prevention techniques by ISPs as a meaningful step in securing inter-domain routing.

A. Operator Engagement

This study’s preprint was posted to the NANOG and RIPE operator mailing lists in June and August 2020. While operator response was limited, email correspondence around the study did yield some helpful insights. One European AS operator claimed that at least one (and possibly additional) transit networks deployed techniques similar to Peerlock around 2007, roughly ten years before NTT’s codification of the method [18]. Separately, a Tier 1 network operator suggested that 1) differences in network automation sophistication could account for observed uneven filtering within the peering clique and 2) defensive filtering may have partially mitigated the Verizon/Cloudflare leaks detailed in Section II-B. This idea is supported by Cloudflare’s discussion on the incident [23]. Cloudflare identifies some networks that filtered the leaks (including ASes 1299 Telia, 2914 NTT, and 7018 AT&T). Bandwidth graphs presented in that post indicate little or no impact on Cloudflare-to-AT&T data plane operation, while Cloudflare traffic to leak propagator Verizon was drastically reduced for hours after the incident.

B. Future Directions

Widespread adoption of Peerlock will likely depend on addressing scalability issues. Rule configuration currently requires non-standard, manual out-of-band communication between protector/protected ASes. Automating this process is a crucial step in extending Peerlock beyond core networks. Communities designating authorized upstreams for routes, as proposed in [16], could take the place of out-of-band communication. Alternatively, RPKI registration of direct/indirect customers [17] could distribute trusted topological information relevant to filtering.

ACKNOWLEDGEMENTS

The authors would like to thank Samuel Jero for his helpful guidance in the shepherding process. The detailed and thoughtful comments/criticisms from our NDSS reviewers are also appreciated, especially within the context of the ongoing COVID-19 pandemic. Finally, we would like to recognize Job Snijders at NTT Communications for presenting the Peerlock system. This study was supported by the National Science Foundation under Grant No. 1850379.

REFERENCES

- [1] R. Anwar, H. Niaz, D. Choffnes, Í. Cunha, P. Gill, and E. Katz-Bassett, "Investigating interdomain routing policies in the wild," in *Proceedings of the 2015 Internet Measurement Conference*, 2015, pp. 71–77.
- [2] H. Birge-Lee, L. Wang, J. Rexford, and P. Mittal, "Sico: Surgical interception attacks by manipulating bgp communities," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019, pp. 431–448.
- [3] R. Bush, O. Maennel, M. Roughan, and S. Uhlig, "Internet optometry: assessing the broken glasses in internet reachability," in *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement*, 2009, pp. 242–253.
- [4] J. Choi, J. H. Park, P.-c. Cheng, D. Kim, and L. Zhang, "Understanding bgp next-hop diversity," in *2011 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 2011, pp. 846–851.
- [5] M. A. N. for Routing Security, "Manrs filtering guide," <https://www.manrs.org/fisps/guide/filtering/>.
- [6] —, "Manrs site," <https://www.manrs.org/>.
- [7] Gao, Lixin, "On inferring autonomous system relationships in the Internet," *IEEE/ACM ToN*, 2001.
- [8] Y. Gilad, A. Cohen, A. Herzberg, M. Schapira, and H. Shulman, "Are we there yet? on rpki's deployment and security." 2017.
- [9] Y. Gilad, O. Sagga, and S. Goldberg, "Maxlength considered harmful to the rpki," in *Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies*, 2017, pp. 101–107.
- [10] V. Giotsas, M. Luckie, B. Huffaker, and K. Claffy, "Inferring complex as relationships," in *Proceedings of the 2014 Conference on Internet Measurement Conference*, 2014, pp. 23–30.
- [11] Goodin, Dan, "Google goes down after major BGP mishap routes traffic through China," <https://arstechnica.com/information-technology/2018/11/major-bgp-mishap-takes-down-google-as-traffic-improperly-travels-to-china/>, 2018.
- [12] T. Hlavacek, I. Cunha, Y. Gilad, A. Herzberg, E. Katz-Bassett, M. Schapira, and H. Shulman, "Disco: Sidestepping rpki's deployment barriers," in *NDSS*, 2019.
- [13] G. Huston and G. Michaelson, "Rfc 6483 - route origin validation," 2012.
- [14] IETF, "IRR RPSL Reference Guide," <https://tools.ietf.org/html/rfc2622>.
- [15] M. N. Inc, "List of routing registries," <http://www.irr.net/docs/list.html>.
- [16] J. H. J. Snijders, "BGP Route Leak Protection Community," <https://www.ietf.org/archive/id/draft-heitz-idr-route-leak-community-00.txt>, 2017.
- [17] M. S. J. Snijders and M. Aelmans, "RPKI Autonomous Systems Cones," <https://tools.ietf.org/html/draft-ietf-grow-rpki-as-cones-02>, 2020.
- [18] Job Snijders, "NTT Peer Locking," http://instituut.net/~job/peerlock_manual.pdf, 2016.
- [19] E. Katz-Bassett, D. R. Choffnes, Í. Cunha, C. Scott, T. Anderson, and A. Krishnamurthy, "Machiavellian routing: improving internet availability with bgp poisoning," in *Proceedings of the 10th ACM Workshop on Hot Topics in Networks*, 2011, pp. 1–6.
- [20] E. Katz-Bassett, C. Scott, D. R. Choffnes, Í. Cunha, V. Valancius, N. Feamster, H. V. Madhyastha, T. Anderson, and A. Krishnamurthy, "Lifeguard: Practical repair of persistent route failures," *ACM SIGCOMM Computer Communication Review*, vol. 42, no. 4, pp. 395–406, 2012.
- [21] Kephart, Nick, "Finding and Diagnosing BGP route leaks," <https://blog.thousandeyes.com/finding-and-diagnosing-bgp-route-leaks/>, 2015.
- [22] B. Kuerbis and M. Mueller, "Internet routing registries, data governance, and security," *Journal of Cyber Policy*, vol. 2, no. 1, pp. 64–81, 2017.
- [23] Levy, Martin, "The deep-dive into how Verizon and a BGP Optimizer Knocked Large Parts of the Internet Offline Monday," <https://blog.cloudflare.com/the-deep-dive-into-how-verizon-and-a-bgp-optimizer-knocked-large-parts-of-the-internet-offline-monday/>, 2019.
- [24] M. Lepinski and K. Sriram, "RFC 8205 - BGPsec protocol specification," *IETF*, 2013.
- [25] M. Lepinski and S. Kent, "RFC 6480 - An Infrastructure to Support Secure Internet Routing," <https://tools.ietf.org/html/rfc6480>, 2012.
- [26] Madory, Doug, "Why Far-Flung Parts of the Internet Broke Today," <https://dyn.com/blog/why-the-internet-broke-today/>, 2014.
- [27] —, "Widespread impact caused by Level 3 BGP route leak," <https://blogs.oracle.com/internetintelligence/widespread-impact-caused-by-level-3-bgp-route-leak>, 2017.
- [28] T. McDaniel, J. M. Smith, and M. Schuchard, "The maestro attack: Orchestrating malicious flows with bgp," in *Proceedings of the 16th EAI International Conference on Security and Privacy in Communication Networks (SecureComm)*, 2020.
- [29] R. NCC, "Routing information service," <https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris>.
- [30] —, "Ripe ncc irr database non-authoritative route object clean-up," <https://www.ripe.net/publications/docs/ripe-731>, 2019.
- [31] NIST, "RPKI Deployment Monitor," <https://rpki-monitor.antd.nist.gov/>, 2020.
- [32] U. of Oregon, "Routeviews project," <http://www.routeviews.org>.
- [33] Oliveira, Ricardo and Pei, Dan and Willinger, Walter and Zhang, Beichuan and Zhang, Lixia, "The (in) completeness of the observed internet AS-level structure," *IEEE/ACM ToN*, 2010.
- [34] C. Orsini, A. King, D. Giordano, V. Giotsas, and A. Dainotti, "Bgpstream: a software framework for live and historical bgp data analysis," in *IMC*, 2016, pp. 429–444.
- [35] J. Schlamp, R. Holz, Q. Jacquemart, G. Carle, and E. W. Biersack, "Heap: reliable assessment of bgp hijacking attacks," *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 6, pp. 1849–1861, 2016.
- [36] B. Schlinker, T. Arnold, Í. Cunha, and E. Katz-Bassett, "Peering: Virtualizing bgp at the edge for research," in *Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies*, 2019, pp. 51–67.
- [37] Schuchard, Max and Geddes, John and Thompson, Christopher and Hopper, Nicholas, "Routing Around Decoys," in *ACM CCS*, 2012.
- [38] Schuchard, Max and Mohaisen, Abdelaziz and Foo Kune, Denis and Hopper, Nicholas and Kim, Yongdae and Vasserman, Eugene Y, "Losing control of the internet," in *ACM CCS*, 2010.
- [39] P. Sermpezis, V. Kotronis, P. Gigis, X. Dimitropoulos, D. Cicalese, A. King, and A. Dainotti, "Artemis: Neutralizing bgp hijacking within a minute," *IEEE/ACM Transactions on Networking*, vol. 26, no. 6, pp. 2471–2486, 2018.
- [40] Shapelez, Alex, "This is how you deal with route leaks," <https://habr.com/en/company/qtrator/blog/495260/>, 2020.
- [41] J. M. Smith, K. Birkeland, T. McDaniel, and M. Schuchard, "Withdrawing the bgp re-routing curtain: Understanding the security impact of bgp poisoning through real-world measurements," in *NDSS*, 2020.
- [42] J. M. Smith and M. Schuchard, "Routing around congestion: Defeating ddos attacks and adverse network conditions via reactive bgp routing," in *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2018, pp. 599–617.
- [43] Snijders, J., Heasley, J. and Schmidt, M., "Use of bgp large communities," <https://tools.ietf.org/html/rfc8195>, 2017.

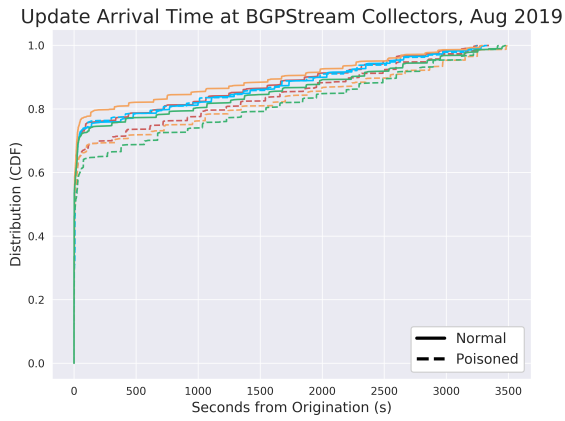
- [44] Snijders, Job, “Everyday Practical BGP Filtering,” <https://peerlock.net>, 2016.
- [45] K. Sriram, D. Montgomery, D. McPherson, E. Osterweil, and B. Dickson, “Rfc 7908 - problem definition and classification of bgp route leaks,” <https://tools.ietf.org/html/rfc7908>.
- [46] Sriram, K. and Azimov, A., “Methods for detection and mitigation of bgp route leaks,” <https://tools.ietf.org/pdf/draft-ietf-grow-route-leak-detection-mitigation-02.pdf>, 2020.
- [47] Sriram, Kotikalapudi and Montgomery, Douglas C., “Resilient Interdomain Traffic Exchange: BGP Security and DDoS Mitigation,” *NIST*, 2019.
- [48] Strickx, Tom, “How Verizon and a BGP Optimizer Knocked Large Parts of the Internet Offline Today,” <https://blog.cloudflare.com/how-verizon-and-a-bgp-optimizer-knocked-large-parts-of-the-internet-offline-today/>, 2019.
- [49] Tran, Muoi and Kang, Min Suk and Hsiao, Hsu-Chun and Chiang, Wei-Hsuan and Tung, Shu-Po and Wang, Yu-Su, “On the Feasibility of Rerouting-based DDoS Defenses,” in *IEEE S&P*, 2019.
- [50] UCSD-CAIDA, “CAIDA AS Ranking,” <https://asrank.caida.org/>, 2020.
- [51] —, “CAIDA AS Relationship dataset,” <https://www.caida.org/data/as-relationships/>, 2020.
- [52] M. Wählisch, O. Maennel, and T. C. Schmidt, “Towards detecting bgp route hijacking using the rpki,” *ACM SIGCOMM Computer Communication Review*, vol. 42, no. 4, pp. 103–104, 2012.
- [53] H. Yan, R. Oliveira, K. Burnett, D. Matthews, L. Zhang, and D. Massey, “Bgpmon: A real-time, scalable, extensible monitoring system,” in *2009 Cybersecurity Applications & Technology Conference for Homeland Security*. IEEE, 2009, pp. 212–223.
- [54] Zmijewski, Earl, “Indonesia Hijacks the World,” <https://dyn.com/blog/indonesia-hijacks-world/>, 2014.

APPENDIX

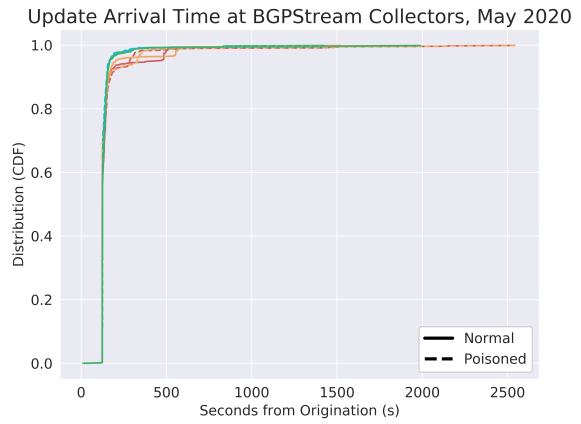
Update Propagation: Before (August 2019) and after (May 2020) conducting our control-plane experiments in Section IV, we performed simple tests to measure 1) the time distribution for BGP update arrivals at RIPE/RouteViews collectors for normal and poisoned advertisements issued from PEERING, and 2) the time distribution for unique ASes seen on AS PATHs in those updates. The latter is most critical for our experiments, as we build our filtering inferences from the presence/absence of ASes on update AS PATHs.

These tests consisted of an explicit /24 withdrawal followed by a one hour waiting period, then a normal /24 advertisement. We listened for updates for the /24 at all BGPStream collectors, and recorded the arrival times of updates for the advertised prefix for one hour. We also noted when unique ASes were first seen on the updates’ AS PATHs. This process (withdraw, update, listen) was repeated five times. We conducted the same process with a poisoned /24 advertisement, for a total of 10 advertisements per experiment.

The results are shown below, Figs.12 and 13. About 80% of updates triggered by a normal or poisoned /24 advertisement that arrived within an hour were received within 30 minutes post-origination in the August experiment. In May, more than 95% of updates fell within this period. More importantly, for every experiment, all unique ASes seen on update paths over the hour listening window arrived within the first 25 minutes post-origination. Over 95% of unique ASes were seen within 7 minutes post-origination.

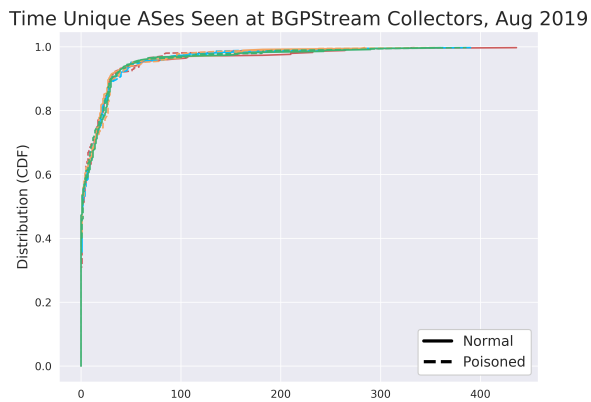


(a) August results.

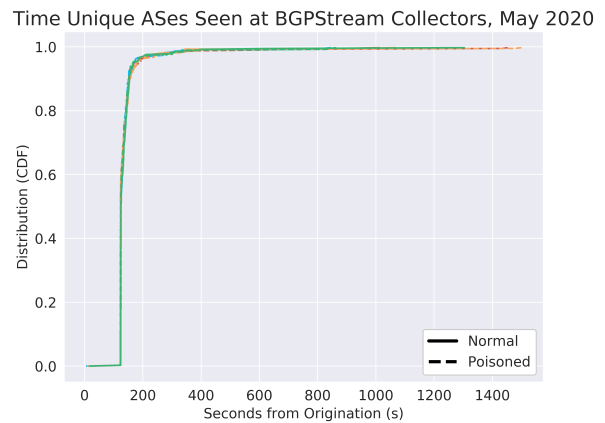


(b) May results.

Fig. 12: Update arrival time CDF. Each of five propagation experiments is illustrated in a different color.



(a) August results.



(b) May results.

Fig. 13: Unique AS arrival time CDF. Each of five propagation experiments is illustrated in a different color.