Project Summary: **TWC: TTP Option: Small: Collaborative: Detecting and Characterizing Internet Traffic Interception based on BGP Hijacking**

Recent reports have highlighted incidents of massive Internet traffic interception executed by rerouting BGP paths across the globe (affecting banks, governments, entire network service providers, etc.). The potential impact of these attacks can range from massive eavesdropping to identity spoofing or selective content modification. In addition, executing such attacks does not require access or proximity to the affected links and networks, posing increasing risks to national security. Worse yet, the impact of traffic interception on the Internet is practically unknown, with even large-scale and long-lasting events apparently going *unnoticed* by the victims.

As reported by Renesys Corporation in November of last year, there is evidence that traffic interception events are growing more frequent, but there are no validated methods to immediately detect them or evaluate their impact. The architectural innovation that mitigates this inherent protocol design flaw exploited by such attacks, is slow to take off, suggesting that this vulnerability will persist, leaving our critical communication infrastructure exposed.

Because of their complex dynamics, and the number of different actors involved on a global scale, devising effective methodologies for the detection and characterization of traffic interception events requires empirical and timely data (e.g., acquired while the event is still ongoing). Such data must be a combination of passive BGP measurements and active measurements (such as traceroutes), since the mechanism triggering the attack operates on the inter-domain routing control plane, but the actual impact is only verifiable in the data plane. By leveraging our experience in measuring and investigating events affecting inter-domain communication and leveraging our measurement and data processing infrastructure, we propose to: (i) investigate, develop, and experimentally evaluate novel methodologies to automatically detect traffic interception events and to characterize their extent, frequency, and impact; (ii) extend our measurement infrastructure to detect in near-realtime and report episodes of traffic interception based on BGP hijacking; (iii) document such events, providing datasets to researchers as well as informing operators, emergency-response teams, law-enforcement agencies, and policy makers. In characterizing their impact, we will quantify increased latency along observed paths, the magnitude of the incident in terms of number of ASes and prefixes intercepted, and the social/political implications of interceptions that take traffic across national borders. We will augment our active measurement framework with algorithmic simulations of BGP routing policies, and qualitative analysis of the organizations involved, to better understand the both technical and political effects of hijacks.

The **intellectual merit** of our proposal lies in our proposed approach to developing scientific methods that can detect and characterize the impact of traffic interception attacks, as well as technology and infrastructure that can demonstrate the value of the developed methods. The results of this project will include efficient techniques for early detection of such events, providing a foundation for future development of reaction and mitigation strategies, and enabling more rigorous pursuit of cybersecurity research.

Broader impacts. Our goal is to advance Internet infrastructure security by creating novel methodologies and instrumentation, as well as improving our understanding of phenomena such as traffic interception. We will inform operators as well as law-enforcement agencies and policy makers with timely empirical data. We will engage faculty, a postdoc, and a graduate student in our project activities, and potentially create collaborations with universities to provide experimental use of our tools and data, creating an immediate link between research and education. We will disseminate project results via conferences, web sites, archived video lectures and blogs.

Keywords: Traffic Interception; BGP; Hijacks; Network Measurement; Internet Monitoring;

Project Description: TWC: TTP Option: Small: Collaborative: Detecting and Characterizing Internet Traffic Interception based on BGP Hijacking

1 Motivation and goals

The Border Gateway Protocol (BGP) is the protocol used to route traffic between autonomous systems (ASes) on the Internet. Designed when the Internet was comprised of a few cooperative ASes, BGP lacks any form of path or origin validation, leaving it extremely vulnerable to attacks and misconfiguration. One example is the fact that networks can advertise illegitimate paths that redirect traffic destined for another network to themselves – known as BGP *hijacking* [1]. Researchers, operators and media have documented and studied BGP hijacks that impact network reachability. Such events either create a traffic *black hole* (*e.g.,* because of a route leak, or to perform a denial-of-service attack) or illicitly use the victim's address block, *e.g.*, to execute an anonymized spamming campaign, or otherwise impersonate the victim [2–5]. However, in 2010, China Telecom's hijack of traffic destined to 50,000 prefixes demonstrated that large-scale traffic *interception* (*i.e.,* where hijacked traffic eventually reaches its intended destination) can also occur on the Internet [3, 6]. While this incident gained wide press exposure [7] and attention from the U.S. government [8], it was largely assumed such interception incidents were usually unintentional and in any event too rare to merit concerted attention.

In November 2013, Renesys Corporation [9] documented and brought to public attention several incidents of massive Internet traffic interception resulting in traffic detouring through unintended countries and even continents (affecting financial institutions, governments, and VoIP providers), and showed that such behavior is increasing. The potential harm caused by these large-scale "man-in-the-middle" attacks can be nefarious, ranging from massive eavesdropping to identity spoofing or selective content modification. In addition, executing such attacks does not require access or proximity to the affected links and networks, posing increasing risks to national security. Worse yet, the impact of traffic interception on the Internet is practically unknown, with even large-scale and long-lasting events apparently often going *unnoticed* by the victims [9]. The natural fix to such design issue is the adoption of BGPSEC and RPKI [10, 11]. However, despite standardization efforts [12] deployment of these technologies is slow [13], and even these measures will not completely eliminate the threat of traffic interception on the Internet [14]. This vulnerability will persist, leaving our critical communication infrastructure dramatically exposed.

Characterizing and mitigating traffic interception faces several challenges that distinguish it from other well-studied types of prefix hijacks. First, unlike prefix hijacks that drop traffic or impersonate the victim AS, the impact of interception is much more subtle, since traffic can still reach the intended destination. Second, detecting and confirming these incidents requires combined analysis of path changes in the control plane observed via BGP data as well as data-plane measurements of traffic paths and their associated delays, since the mechanism triggering the attack operates on the inter-domain routing plane, but the impact is only measurable in the data plane *during the event*. This lack of ongoing targeted measurement of hijacks has hindered attempts to characterize or quantify their impact [6]. Finally, any method for identifying interceptions needs to be efficient (due to the potentially large number of paths and network destinations being monitored) and highly accurate (due to the political implications of interceptions that cross national boundaries [8]).

We propose to address these challenges by devising a methodology specifically designed for this type of attack, based on agile analysis of control and data plane measurement (Section 3.2). This methodology will be supported by a measurement infrastructure that we will extend to per-

form targeted active measurements when suspicious events occur (Section 3.1). Efficient correlation of diverse data sources (*e.g.,* BGP live feeds, AS relationships, targeted active probing, historical traceroute data) will enable accurate detection, characterization and quantification of impact. We will validate and refine our methods through collaboration with operators and other researchers. Finally, we will use our methodology and infrastructure to produce and disseminate data about traffic interception on the Internet to relevant communities: ISPs, Internet security researchers, standards bodies, and policy makers (Section 3.3).

2 Background and Related Work

We present background on BGP's vulnerability to hijacks and efforts to prevent and detect them.

2.1 Insecurity of BGP

BGP was designed when the Internet was comprised of a few ASes with strong social and institutional incentives to cooperate. The original underlying architecture included no mechanisms to validate or authenticate messages sent by ASes, which makes it is highly vulnerable to routing incidents caused by misconfiguration or attacks [3, 4, 15], including hijacking. These incidents include large-scale *route leaks* [16], where an AS originates a large number of prefixes allocated to other ASes (*e.g.,* the China Telecom incident [3]) and more suspicious forms of *path manipulation*, where an AS may announce a path that does not actually exist in the AS-graph [14]. In addition to these two vulnerabilities – which are prevented by RPKI [11] and BGPSEC [10], respectively – there are two more techniques ASes may use to intercept traffic which are not solved by these protocols. The first is interception by an AS that *removes AS path prepending* by the victim AS [17]. AS path prepending is when an AS announces a BGP path with its AS number repeated multiple times, to make the path through a given neighbor less desirable. By removing AS prepending, an intercepting AS can make their path to the victim appear more attractive, and thus intercept a larger volume of traffic. The second type of interception, which is not resolved by RPKI and BGPSEC, is interception where an AS is willing to advertise a path for which it does not obtain revenue [14]. These paths are said to *violate the valley-free assumption*, which states that an AS will advertise paths through its customers to all neighboring ASes, but will only advertise paths through its peers and providers to customers. While announcements violating the valley-free assumption are not technically attacks, since an AS may legitimately choose to transit traffic for which it does not obtain revenue, we plan to monitor instances of these violations for ASes who may use these announcements for interception.

2.2 Resource Public Key Infrastructure and BGPSEC

There have been extensive efforts in the standards bodies to design solutions to overcome BGP's vulnerabilities [12]. The most recent and promising are the Resource Public Key Infrastructure (RPKI) [18, 19] and BGPSEC [10]. RPKI is a hierarchical public key infrastructure designed to certify ownership of IP prefixes by ASes, currently being standardized by the IETF [11, 12] and implemented by the regional Internet registries (RIRs) [18, 19]. Using the RPKI and a chain of valid certificates from a trust anchor, network operators can determine if the AS originating an IP prefix has been authorized to do so. BGPSEC [10] uses the RPKI to manage cryptographic keys which are used by ASes deploying the protocol to sign and validate BGP routing announcements: an AS that receives a signed announcement is able to validate that every AS on the path actually sent the corresponding announcement.

While RPKI and BGPSEC would alleviate many of the security problems that currently plague the interdomain routing system, their deployment faces many political, technical, and economic challenges such as agreement on a trust anchor for RPKI [20], correct configuration of RPKI [21] by ISPs, and lack of incentives for networks to deploy these protocols [13]. Additionally, RPKI and BGPSEC only ensure that announcements are valid, they do not prevent against ASes intentionally violating the valley-free policy and exporting valid paths to their peers and providers in order to intercept traffic [14]

2.3 Detecting BGP hijacks

While efforts to prevent BGP hijacks require massive industry coordination, systems that detect these incidents can be implemented by a single enterprise (*e.g.,* Renesys [22]) or even an academic institution (*e.g.,* Cyclops [23]). As a result, many research efforts focus on detection and characterization of such attacks, proposing approaches based on the analysis of control-plane [16, 24–27], data-plane information [28, 29], as well as their combination [30, 31].

Detection methods based on control-plane measurements typically rely on detection of Multiple Origin AS (MOAS) conflicts, where a single prefix is originated by multiple ASes. Once MOAS conflicts are detected, these methods filter false positives, using information provided by the owner of the prefix [27] or additional data sources [16] such as the list of providers of the potential-victim network. Distinguishing routine traffic engineering from attacks or misconfigurations using only control plane (BGP) data is challenging; most techniques yield many false positives or only identify very large events with confidence [16]. Approaches that monitor data-plane connectivity [29] can detect a drop in traffic reaching the victim, but in order to distinguish a hijack from another network failure they usually adopt prefix-owner-based solutions. Zhang *et al.* , for example, generate an alert whenever the reachability of a monitored prefix from a set of topologically close providers is hindered [29]. Similarly, Zheng *et al.* try to detect prefix hijacks by identifying significant path changes in traceroutes to a prefix from different vantage points [28]. These approaches can be faster than those based on existing BGP feeds, which are limited by their update period (*e.g.,* RIPE releases BGP updates every 15 minutes [32]). However, they face scalability challenges and usually can only be deployed per-AS, which precludes global analysis.

Limitations of data-plane and control-plane measurements in isolation led to the development of hybrid approaches, such as using anomalies observed in the control-plane to trigger data-plane measurements from a set of vantage points (*e.g.,* PlanetLab [33] or public looking glass servers [30, 31]). Such approaches look for conflicting data-plane information when a prefix is announced by multiple networks [30] or correlate failed data-plane measurements (*e.g.,* ping) with control-plane information (routing tables) from the same vantage point [31]. While such hybrid approaches can dramatically improve accuracy in hijack detection over methods relying on control-plane or data-plane information, none of the existing methods has been designed explicitly to detect traffic interception. In fact, techniques that rely on failed data-plane reachability as a signature of an incident (*e.g.,* Argus [31], iSPY [29]) preclude identifying traffic interceptions since they do not exhibit this symptom. systems that are able to identify some hijacking attacks leading to a potential interception (*e.g.,* [30] are not well suited to detect subtle interceptions (*e.g.,* those caused by export policy violations [14] or manipulations of path prepending [17]).

2.4 Analysis of Traffic Interception

A review of the current literature highlights a scarcity of research on the detection and analysis of actual traffic interception events. Most prior work focuses on simulating interceptions to under-

stand their potential impact and characterize the risk posed by different ASes [14,34], or performs controlled interceptions to demonstrate their feasibility [2, 34]. The few in-depth studies of interception in the wild have either relied on proprietary data [3, 9] or publicly available datasets that not not amenable to precise quantification of the scale and impact of interception events [6]. To support its proprietary data collection, analysis, and reporting, Renesys maintains hundreds of BGP monitors and vantage points that (per their web page) send on the order of billions of measurements daily [9], but they do not share their data or methods.

PI Gill demonstrated the value of large-scale, real-time data collection in her case study of China Telecom's interception of 50,000 IP prefixes in 2010 [6], which unlike other coverage [3, 35] included publishing a detailed methodology for analyzing its impact. PI Gill's work relied solely on measurement infrastructures available to academic researchers: BGP monitors [32, 36] and ongoing but untargeted data-plane measurements from iPlane [37].

Our proposal will fill gaps in the existing research literature. To support systematic evaluation of new mechanisms we develop to detect hijacks, we will develop infrastructure that can implement and deploy the mechanisms as we design them. This infrastructure will allow us to not only track interception events in real-time, but also to quantify their impact. We will publish our techniques and analysis results for evaluation and reproducibility by others, and to promote a more complete understanding of the prevalence and impact of role of BGP hijacking and traffic interception events on the Internet.

3 Proposed Work

Our proposed work leverages our existing infrastructure and extends it to enable near real-time detection of traffic interception events (Section 3.1). Since continuous distributed active measurements for each routed network prefix is not feasible due to measurement overhead, we take a two-step approach. First, we continuously monitor the control-plane for anomalies (Section 3.2.1) and then use on-demand data-plane measurements to diagnose interception (Section 3.2.2) and quantify its impact (Section 3.2.3). Finally, we will disseminate data from our platform and gain feedback from operators, researchers, and policymakers. (Section 3.3).

3.1 Task 1: Extend and refine infrastructure for data collection and analysis

Figure 1 illustrates the main components of our existing and proposed infrastructure for data collection and analysis of interception activity. Existing software and hardware components of the architecture are colored blue. We propose to develop and integrate the components colored red in the figure, whose functionality we describe in more detail in Section 3.2.

1. **Collection and preprocessing of BGP feeds.** We will process BGP measurements collected by three distinct BGP data collection projects that publish at different latencies:
   - Routeviews [36]: 329 peers; typically less than 20 minutes latency;
   - RIPE RIS [32]: 209 peers; latency varies between 10 minutes and 2 hours;
   - BGPmon [38]: 60 peers (almost entirely overlapping with Routeviews); live stream of BGP updates, few seconds of latency.

   We (CAIDA) have deployed functionality to continuously collect and process BGP feeds from these three sources (blue box on left of Figure 1) as part of an existing SaTC project to detect large-scale Internet outages ("SaTC:TTP Detection and Analysis of Large-Scale Internet Infrastructure Outages", CNS-1228994, Sep 2012 - Aug 2015, Claffy and Dainotti) [39]. We also developed and deployed a software framework (*libbgpstream*) [39] to transparently collect and simultaneously process heterogeneous BGP feeds, and provide a uniform layer

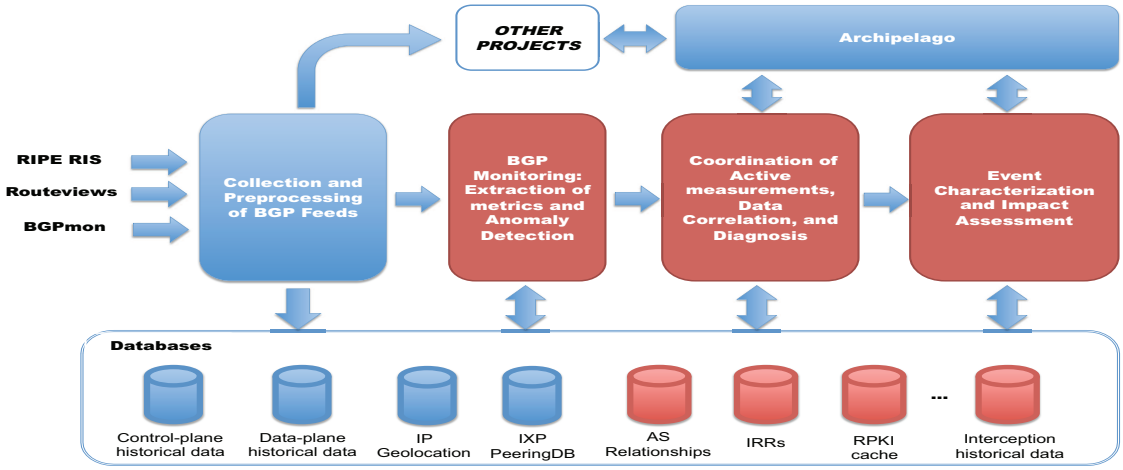

Figure 1: *Our infrastructure for data collection and analysis. Portions in red denote components to be developed and integrated as part of this proposal, whereas components deployed within previous or ongoing synergistic projects are in blue.*

Figure 2: *As of November 2013, there are 83 Ark monitors in 36 countries.*

to access the data. This flexibility will help us evaluate how the trade-off between BGP monitoring latency and coverage impact our methodology.

2. **BGP monitoring: extraction of metrics and anomaly detection.** In this task, we will add software modules to (i) extract specific metrics from BGP feeds (*e.g.*, number of prefixes advertised per AS) and (ii) monitor them for anomalies that may be related to suspicious events (Section 3.2.1).
3. **Archipelago (Ark) active measurement infrastructure.** Ark [40] is CAIDA's active measurement infrastructure [41] running software that enables it to operate many active monitors as a coordinated secure measurement platform. Figure 2 depicts the 83 Ark monitors currently deployed. Ark supports rapid prototyping by promoting software development at a high-level of abstraction using dynamic scripting languages and pre-built APIs and services. It uses Ruby [42] for measurements and related libraries, and to control measurement modules, including *scamper* [43], a flexible active measurement tool supporting ping and several traceroute variants (TCP, UDP, and ICMP, and Paris [44]). Ark also supports decentralized measurement processes executing autonomously at each monitor and communicating as needed, for example, to trigger further measurements or analyses based on observed events.

We will use Ark's APIs and services to implement logic for conducting targeted active measurements when suspicious events occur. We will also coordinate with other Ark's funded activities in order to deploy additional Ark monitors in locations that are beneficial to our methodology. For example, Ark monitors located within ASes that peer with either Routeviews or RIPE RIS can support systematic comparison of data-plane (traceroute-measured) AS paths with (control-plane) BGP AS paths. (Ark currently has 26 monitors located within ASes that peer with either Routeviews or RIPE RIS).

4. **Supporting Databases.** Section 3.2 describes how several sources of Internet measurement data and meta-data support our methodology, some of which are already integrated into our platform, *e.g.,* IP Geolocation, BGP-prefix-to-AS mapping. To support detection and analysis of suspicious events, we will integrate several additional databases, including: RPKI information from ISPs, AS business relationships inferences, a sliding window of historical traceroute data from Ark, and Internet Routing Registries (IRR) information.
5. **Coordination of active measurements, data correlation, and diagnosis.** This software module takes as input anomalies detected at the control plane, and orchestrates targeted active measurements to capture dynamics of the anomaly. It will then try to diagnose the event by correlating data from targeted active measurements with control-plane measurements (*e.g.,* AS paths inferred from traceroutes vs. AS paths announced via BGP) and with other information in our supporting databases, such as previous traceroutes towards the same prefixes (Section 3.2.2).
6. **Event characterization and impact assessment.** When an event is classified as interception, this software module performs additional measurements, runs simulations of the AS-level topology based on AS paths advertised on the control plane, and combines this information with supporting data or meta-data to quantify its impact, *e.g.,* in terms of ASes and geographical regions involved (Section 3.2.3).

3.2 Task 2: Design a method for detecting and characterizing traffic interception

We will use the infrastructure described in Section 3.1 to detect interceptions by coupling passive monitoring of BGP messages (Section 3.2.1) with targeted active measurements (Section 3.2.2), to provide near real-time detection of traffic interception events on the Internet. We will also combine empirical data, simulations, and qualitative analysis to understand the technical and political impact of these incidents (Section 3.2.3).

3.2.1 BGP measurements and analysis

We will use the interface to BGP data provided by our infrastructure to discover anomalies including but not limited to the following:

1. **Multiple Origin AS (MOAS).** When an AS that has not previously originated a prefix begins to do so, we will flag it as an anomaly. However, there are many legitimate reasons that a prefix may be originated by multiple ASes (*e.g.,* anycast, IP transfers, prefix de-aggregation, sibling-ASes [45]). We will use Internet Routing Registries data (IRRs) [46], peeringdb [47], and inferred AS relationships [48] in order to filter out MOAS conflicts that appear legitimate. For example, if IRRs confirm that two ASes are part of the same organization, we do not consider it suspicious for one AS to announce prefixes belonging to the other. Similarly, it is not anomalous for ASes that are members of an IXP to announce a prefix belonging to the IXP's AS [16]. Using these datasets and information collected from the operator community (*e.g.,* via the web dashboard described in Section 3.3) we will curate a list of MOAS

violations that are a normal part of the Internet's routing system and thus do not require further analysis through active measurements.

2. **Increase in prefixes originated by each AS.** In the case of the China Telecom interception, the incident resulted from a leak of 50,000 prefixes by China Telecom's network [3], observable as 50,000 new prefixes originated by China Telecom's data center AS 23724. This rapid increase in prefixes announced by an AS is a strong indicator of a large-scale incident *e.g.,* a route leak [16]. To decrease our false negative rate, we will experiment with varying thresholds on the number of new prefixes an AS announces, and the time period over which they make these announcements.
3. **Violations of the "valley-free" assumption [49].** We will leverage CAIDA's AS topology and relationship inference data [48] to identify BGP paths that violate the valley-free assumption (illustrated in Figure 3b). These violations may occur for many reasons, such as accidental route leaks, intentional agreements between ASes, or even as a result of errors in our AS relationship inferences. We will consider violations of the valley-free assumption as anomalies to further investigate with active measurements. However, as with MOAS anomalies, we will work with providers to curate a list of known exceptions to this assumption, which will not only reduce the false positives at this stage, but may have the added benefit of helping to improve the accuracy of AS relationship inferences.
4. **New edges in the AS graph.** We will monitor the set of edges observed in BGP announcements (similar to how Argus tracks pairs of ASes observed in BGP paths [31]). Using only control-plane data it is incredibly challenging to differentiate new edges that appear for legitimate reasons from path poisoning (Figure 3c). Data plane measurements will help us validate that new edges are not due to interception, in which case we will add them to a list of previously seen edges. As failures happen and ASes explore backup paths, over time we should gain a more complete view of all edges we should expect to see.
5. **Inconsistent prepending announcements.** A recently documented interception attack involves manipulating AS prepending in paths [17] (Figure 3d). In this attack, the malicious AS strips path prepending from an AS path prior to forwarding an announcement. For each prefix, we can monitor the amount of prepending the originating AS adds before forwarding an announcement to each of its upstream providers, and flag sudden changes as anomalies.

While these criteria may seem intuitive, Internet routing often violates our mental models of how it should work. Correlation with external data sources such as IRRs [46], and peeringdb [47] as well as feedback from operators will be critical in validating and refining our criteria. Even once our criteria are refined, implementing these criteria on a system that runs in real time is non-trivial. We will leverage CAIDA's experience building measurement and monitoring systems and PI Gill's experience designing algorithms that model routing policies to develop technology to support a system that provides timely validation and notification of interception incidents.

3.2.2 Diagnosing BGP interception with targeted active measurements

For each event detected as suspicious, the module described in the previous section generates relevant meta-data, including which type of anomaly was observed, potential victim prefixes and ASes, and relevant AS paths to measure. Based on this meta-data, we will execute traceroutes from all Ark monitors to potential victim prefixes, and convert the resulting IP-level paths into AS-level paths using state-of-the-art techniques (*e.g.,* [50, 51]). When possible, we will use historical traceroute data from Ark in order to select IP addresses within each prefix that have previously responded to traceroute. (Our measurements simply require that the traceroute can reach the given prefix not a specific host.) We will also leverage Ark's ability to execute UDP- or

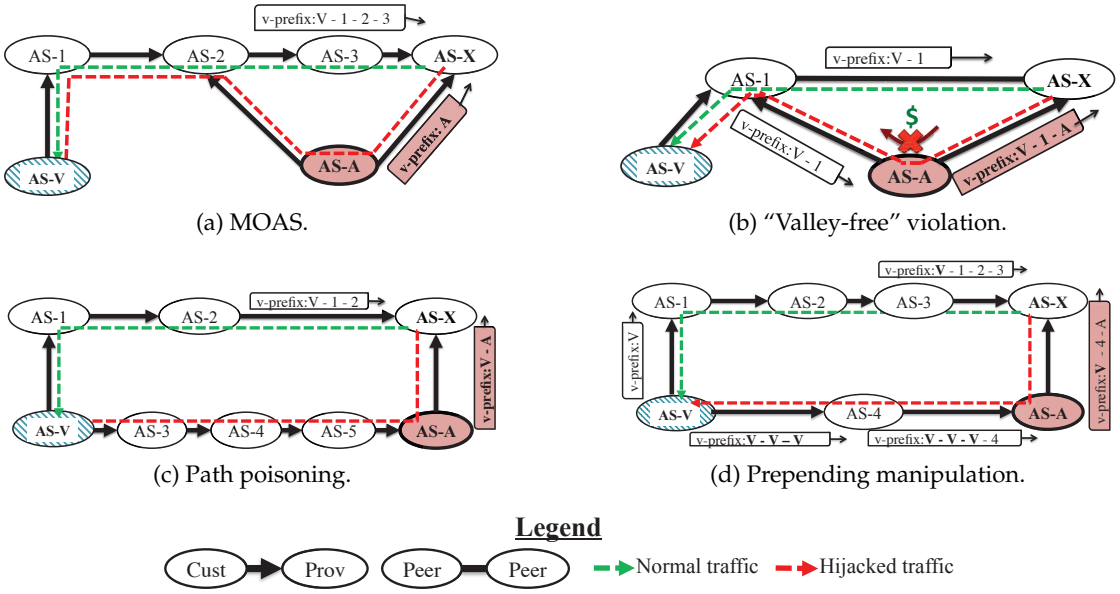

Figure 3: *Examples of interceptions that may be detected using the indicators in Section 3.2.1*

TCP- traceroutes, which can override filtering of ICMP enforced by some providers. We will compare these traceroute-inferred with BGP-inferred paths, and map the results of this comparison to different scenarios in Table 1. In the table, rows represent different anomalies observable in the control-plane (BGP) (Section 3.2.1), and the columns depict possible outcomes of the comparison: AS paths match, do not match, or active probing did not reach the targeted prefix. In practice, we will need to interpret Table 1 separately for each vantage point, since not all vantage points may be impacted by the interception. Corresponding results from multiple vantage points will reflect portions of the AS topology also affected by the event. Table 1 depicts the four types of interception illustrated in Figure 3.

Interception 1: Multiple origin interception (Figure 3a). This type of interception results from an AS originating prefixes not allocated to it in the control-plane (*e.g.*, MOAS violation or an AS announcing an unusually large number of prefixes). If the data-plane path to the prefix is not consistent with the path in the control-plane, and it traverses the potential attacker AS, we confirm that it is a case of interception.

Interception 2: "Valley-free" violation (Figure 3b). This type of interception results in violations of the valley-free policy in the control plane. We use data-plane measurements to confirm that traffic is indeed flowing over the path announced in the control-plane. This type of interception is more subtle since it does not actually violate the validity of the announced path (*i.e.*, the control- and data-plane paths will match). Thus, when we observe these types of symptoms we will reach out to relevant victims to verify that these announcement are actually unintended and not pre-arranged transit swaps. In the absence of individuals who can validate these incidents, we will perform temporal correlations to gain more insight into this type of event.

Interception 3: Path poisoning interception (Figure 3c). This type of interception is characterized by a previously unseen edge appearing in control-plane messages. Similar to MOAS-based interceptions, if the data-plane path differs from the control-plane path we conclude that this incident is the result of an interception.

Interception 4: Prepending manipulation (Figure 3d). Interception that relies on manipulating

Table 1: *Interpretation of control-plane and data-plane measurements to diagnose interception (interception events are marked in bold italic fonts). For each event detected as suspicious by the BGP monitoring module, we will execute traceroutes from all Ark monitors to potential victim prefixes, and convert the resulting IP-level paths into AS-level paths using state-of-the-art techniques (*e.g., *[50, 51]). We will compare these traceroute-inferred with BGP-inferred paths, and map the results of this comparison to the different scenarios in the table. Rows represent different anomalies observable in the control-plane (BGP) (Section 3.2.1), and the columns depict possible outcomes of the comparison: AS paths match, do not match, or active probing did not reach the targeted prefix.*

Control-plane anomaly	**Data-plane *vs.* Control-plane path**		
	AS paths match	**AS paths are different**	**Could not reach prefix**
1. MOAS **2. Prefix increase by AS**	no interception, impersonation [5]	***(1) multiple-origin interception***	black-hole hijack
3. "Valley-free" violation	special agreements, misconfiguration, ***(2) interception via valley-free violation***	suspicious data plane and control plane mismatch	misconfiguration, black-hole hijack
4. New edge	new connection, impersonation [5]	***(3) interception via path poisoning***	misconfiguration, black-hole hijack
5. Inconsistent prepend	***(4) interception via prepend manipulation***	suspicious data plane and control plane mismatch	misconfiguration

path prepending is more subtle and trickier to detect. The primary evidence for this sort of interception comes from inconsistent prepending announcements observed in the control-plane. In this case, data-plane measurements serve to confirm that traffic is actually being diverted through the network suspected of performing the interception.

While we have a taxonomy and intuition of how to detect different types of interceptions, there are many challenges our research will have to address. Many signatures of interception are similar to signatures of network failures or other types of hijacking (Table 1). Distinguishing among types of events will require experimenting with different techniques. When we find multiple interpretations of an incident, we will contact the ASes involved, *e.g.,* to distinguish misconfiguration and special agreements from interceptions in the case of valley-free violations.

3.2.3 Using empirical data and simulations to study the impact of interceptions

Detecting interceptions is only one piece of the puzzle in terms of understanding the role they play on the Internet. It is also important to understand their impact in terms of: (1) service impact via increased latency, (2) the magnitude of the incident in terms of number of ASes and prefixes intercepted, (3) duration of the incident and (4) social/political implications of interceptions that take traffic across national borders. We plan to augment our active measurement framework with algorithmic simulations of BGP routing policies, and qualitative analysis of the organizations involved, to better understand both technical and political effects of interceptions.

1. **The impact of interception on network latency.** Traffic interception can result in abnormally high network delays, as traffic detours around the globe. For example, the China Telecom incident resulted in some network latencies increasing by 300ms before traffic reached its intended destination [3]. When we detect an interception activity, we will leverage historical traceroute data gathered by CAIDA's Ark project [40], comparing previous traceroutes between a given source an destination with those executed during the interception event, allowing us to estimate the effect of the interception in terms of increased latency.
2. **Magnitude of interception in terms of ASes and prefixes.** Recognizing the limitations in coverage (number and location of available vantage points) of our measurement infrastructure, we will combine both measurement data and simulation to try to more accurately estimate the number of ASes and prefixes affected by an interception. Using BGP data we will derive an initial count of how many prefixes are impacted by an interception event, and then execute active measurements from Ark nodes to these prefixes to confirm which ASes are *definitely* impacted. We will supplement this data with algorithmic simulations of BGP that use the best available AS-level topology annotated with relationships ([48], developed at CAIDA) and efficient algorithms to compute BGP paths (designed by PI Gill [52]). Since ASes choose paths based on considerations other than shortest path (*e.g.,* cost, geography), cannot use existing shortest path algorithms without modification. Instead, our algorithm performs a three-stage breadth-first search of the AS graph to compute paths chosen by each AS to reach a given destination [52]. We will compute paths to the victim AS as well as the AS performing interception, and compare them. If an AS has a better path (according to a standard model of routing policies [53]) to the intercepting AS than it has to the victim, we conclude that the interception would impact this AS. This analysis will provide an upper bound on the impact of the potential interception and help fill gaps in our active measurements.
3. **Duration of traffic interceptions.** Another important question about the impact of traffic interception is: how long did the interception last? The amount of data potentially intercepted is tied to the duration of the interception, so duration an important metric of impact in the absence of network traces from the impacted networks. When we detect an interception, we will monitor for resolution of the corresponding control-plane anomaly (Section 3.2.1) and follow up with active measurements to verify that the interception is indeed over. Once we have confirmed that the interception is no longer taking place, we can investigate the duration of interception events and perform temporal correlations to observe patterns such as whether some interception events recur frequently.
4. **Social/political impact of interception.** When China Telecom hijacked 50,000 prefixes in 2010, some of which belonged to key US organizations, the impact was not just technical but political [8]. Thus, we will augment our technical measurements of interception with useful meta-data for inferring social and political impact of these incidents. First, we will use data from the regional Internet registries (RIRs) (*e.g.,* ARIN [54], RIPE [55]) to map the

intercepted prefixes to the organizations that registered them, and identify potentially significant targets (*e.g.,* U.S. Patent and Trademark Organization [6] and other prefixes associated with national government organizations). We will also use the RIRs to map which ASes had traffic diverted as part of the hijack to see if any of them may be of political significance. Interception that diverts traffic to a foreign country is likely of more concern than one that diverts it within a single country, so to understand where intercepted traffic flows, we will map our interception data using available IP geolocation information, both from commercial sources [56] and from CAIDA's recent geolocation research project [57].

3.3 Task 3: Collect and disseminate data and knowledge

The proposed project is partly inspired by an ISOC roundtable [58] attended by the PIs and other academics, as well as network operators and policy makers. The roundtable discussions highlighted the need for high-fidelity, unified data and analysis of routing security incidents. We will include the following activities to engage these and other communities, not only to disseminate our results but also to solicit feedback as we try to address the identified identified gap.

1. **Workshop in Year 2.** Early in Year 2 of the project, we will host a 1-2 day workshop at Stony Brook University (SBU) that will include the research team as well as network operators and policy makers (*e.g.,* members of the IETF Secure Interdomain Routing (SIDR) working group and the FCC CSRIC IV working group on BGP Security). The workshop will leverage existing ties between the PIs and these specific groups, as well as the PIs participation in operator forums such as the North American Network Operators Group (NANOG). The workshop will raise awareness of our research and the data it will provide. We will also gather feedback on our detection criteria and progress in Year 1, to help set the course for Year 2.
2. **Maintaining a blog detailing incidents.** The research team will maintain a blog summarizing notable events discovered using the infrastructure. The blog will present detailed analysis of incidents in a format that is understandable, even to non-technical policy makers, and link to more detailed technical reports. The blog will include not only traffic interception analysis, but analysis of other incidents detectable using our infrastructure (*e.g.,* hijacks resulting in traffic black holes and other types of outages) to make it the go-to place for Internet resiliency reporting (similar to Renesys, but transparently sharing data, analysis methods, and software to promote reproducibility).
3. **Engaging the operator community.** The PIs will leverage their prior success engaging the operator community (*e.g.,* [59,60]) to share research results with them. The developed infrastructure will include an interface that vetted operators can use to monitor incidents relating to their networks, or to manually enter any information or corrections that may improve the accuracy of the reported incidents. Once this interface is implemented, the PIs will present the infrastructure at a NANOG meeting, and announce its availability on the NANOG mailing list to bootstrap participation by network operators.
4. **Data access for academics.** The data produced by our infrastructure is useful for performing "what-if" analysis to understand the effectiveness of different routing security deployments (*e.g.,* [14,61]). It is also useful for studies of new detection mechanisms for Internet outages, prefix hijacks, and interceptions. Finally, the data can serve as a starting point for assignments in courses on networking and network security. We will provide access to our data through the CAIDA data repository to maximize its benefit to the academic community.

References

[1] K. Butler, T. Farley, P. McDaniel, and J. Rexford, "A Survey of BGP Security Issues and Solutions," *Proceedings of the IEEE*, vol. 98, no. 1, pp. 100–122, 2010.

[2] A. Pilosov and T. Kapela, "Stealing the Internet: An Internet-Scale Man in the Middle Attack," 2008. Presentation at DefCon 16, http://www.defcon.org/images/defcon-16/dc16-presentations/defcon-16-pilosov-kapela.pdf.

[3] J. Cowie, "Renesys blog: China's 18-minute mystery," 2010. http://www.renesys.com/blog/2010/11/chinas-18-minute-mystery.shtml.

[4] S. Misel, "Wow, AS7007!," 1997. http://www.merit.edu/mail.archives/nanog/1997-04/msg00340.html.

[5] A. Ramachandran and N. Feamster, "Understanding the Network-level Behavior of Spammers," *SIGCOMM Comput. Commun. Rev.*, vol. 36, pp. 291–302, Aug. 2006.

[6] R. Hiran, N. Carlsson, and P. Gill, "Characterizing Large-Scale Routing Anomalies: A Case Study of the China Telecom Incident," in *Passive and Active Measurement* (M. Roughan and R. Chang, eds.), vol. 7799 of *Lecture Notes in Computer Science*, pp. 229–238, Springer Berlin Heidelberg, 2013.

[7] R. McMillan, "A Chinese ISP Momentarily Hijacks the Internet," 2010. `http://www.nytimes.com/external/idg/2010/04/08/08idg-a-chinese-isp-momentarily-hijacks-the-internet-33717.html`.

[8] D. Blumenthal, P. Brookes, R. Cleveland, J. Fiedler, P. Mulloy, W. Reinsch, D. Shea, P. Videnieks, M. Wessel, and L. Wortzel, "Report to Congress of the US-China Economic and Security Review Commission," 2010. `http://www.uscc.gov/annual_report/2010/annual_report_full_10.pdf`.

[9] J. Cowie, "Renesys blog: The New Threat: Targeted Internet Traffic Misdirection," 2013. `http://www.renesys.com/2013/11/mitm-internet-hijacking/`.

[10] R. Austein, S. Bellovin, R. Bush, R. Housley, M. Lepinski, S. Kent, W. Kumari, D. Montgomery, K. Sriram, and S. Weiler, "BGPSEC Protocol Specification," 2012. http://datatracker.ietf.org/doc/draft-ietf-sidr-bgpsec-protocol/.

[11] R. Austein, G. Huston, S. Kent, and M. Lepinski, "Secure Inter-Domain Routing: Manifests for the Resource Public Key Infrastructure." draft-ietf-sidr-rpki-manifests-09.txt, 2010.

[12] The Internet Engineering Task Force, "Secure interdomain routing (SIDR) working group," 2012. http://datatracker.ietf.org/wg/sidr/charter/.

[13] P. Gill, M. Schapira, and S. Goldberg, "Let the Market Drive Deployment: A Strategy for Transitioning to BGP Security," *SIGCOMM Comput. Commun. Rev.*, vol. 41, pp. 14–25, Aug. 2011.

[14] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford, "How secure are secure interdomain routing protocols," *SIGCOMM Comput. Commun. Rev.*, vol. 41, pp. –, Aug. 2010.

[15] M. Brown, "Renesys blog: Pakistan Hijacks YouTube," 2008. http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml.

[16] V. Khare, Q. Ju, and B. Zhang, "Concurrent prefix hijacks: Occurrence and impacts," in *Proceedings of the 2012 ACM Conference on Internet Measurement Conference*, IMC '12, (New York, NY, USA), pp. 29–36, ACM, 2012.

[17] Y. Zhang and M. Pourzandi, "Studying Impacts of Prefix Interception Attack by Exploring BGP AS-PATH Prepending," *2013 IEEE 33rd International Conference on Distributed Computing Systems*, vol. 0, pp. 667–677, 2012.

[18] American Registry for Internet Numbers, "ARIN Resource certification." http://www.arin.net/resources/rpki.html.

[19] RIPE Network Coordination Center, "RIPE NCC Resource Certification." http://www.ripe.net/certification/.
[20] J. Curran, "RPKI and Trust Anchor question," 2013. http://mailman.nanog.org/pipermail/nanog/2013-August/060199.html.
[21] M. Wahlisch, O. Maennel, and T. Schmidt, "Towards Detecting BGP Route Hijacking using the RPKI," in *ACM SIGCOMM Posters*, 2012.
[22] "Renesys." https://www.renesys.com/.
[23] University of California - Los Angeles, "Internet Research Lab – Internet topology collection." http://irl.cs.ucla.edu/topology/.
[24] J. Karlin, S. Forrest, and J. Rexford, "Autonomous Security for Autonomous Systems," *Comput. Netw.*, vol. 52, pp. 2908–2923, Oct. 2008.
[25] C. Kruegel, D. Mutz, W. Robertson, and F. Valeur, "Topology-Based Detection of Anomalous BGP Messages," in *Recent Advances in Intrusion Detection* (G. Vigna, C. Kruegel, and E. Jonsson, eds.), vol. 2820 of *Lecture Notes in Computer Science*, pp. 17–35, Springer Berlin Heidelberg, 2003.
[26] G. Siganos and M. Faloutsos, "Neighborhood Watch for Internet Routing: Can We Improve the Robustness of Internet Routing Today?," in *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*, pp. 1271–1279, 2007.
[27] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang, "PHAS: A Prefix Hijack Alert System," in *Proceedings of the 15th Conference on USENIX Security Symposium - Volume 15*, USENIX-SS'06, (Berkeley, CA, USA), USENIX Association, 2006.
[28] C. Zheng, L. Ji, D. Pei, J. Wang, and P. Francis, "A Light-weight Distributed Scheme for Detecting Ip Prefix Hijacks in Real-time," in *Proceedings of the 2007 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, SIGCOMM '07, (New York, NY, USA), pp. 277–288, ACM, 2007.
[29] Z. Zhang, Y. Zhang, Y. C. Hu, Z. M. Mao, and R. Bush, "iSPY: Detecting IP Prefix Hijacking on My Own," *IEEE/ACM Trans. Netw.*, vol. 18, pp. 1815–1828, Dec. 2010.
[30] X. Hu and Z. M. Mao, "Accurate Real-time Identification of IP Prefix Hijacking," in *Proceedings of the 2007 IEEE Symposium on Security and Privacy*, SP '07, (Washington, DC, USA), pp. 3–17, IEEE Computer Society, 2007.
[31] X. Shi, Y. Xiang, Z. Wang, X. Yin, and J. Wu, "Detecting prefix hijackings in the internet with argus," in *Proceedings of the 2012 ACM Conference on Internet Measurement Conference*, IMC '12, (New York, NY, USA), pp. 15–28, ACM, 2012.
[32] RIPE Network Coordination Center, "RIPE Routing Information Service." http://www.ripe.net/data-tools/stats/ris/routing-information-service.
[33] L. Peterson, T. Anderson, D. Culler, and T. Roscoe, "A blueprint for introducing disruptive technology into the Internet," in *Hotnets*, pp. 59–64, 2002.
[34] H. Ballani, P. Francis, and X. Zhang, "A Study of Prefix Hijacking and Interception in the Internet," *SIGCOMM Comput. Commun. Rev.*, vol. 37, pp. 265–276, Aug. 2007.
[35] BGPMon, "China telecom hijack," 2010. http://bgpmon.net/blog/?p=282.
[36] University of Oregon, "Route views project." http://www.routeviews.org/.
[37] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani, "iPlane: An Information Plane for Distributed Services," in *Proceedings of the 7th Symposium on Operating Systems Design and Implementation*, OSDI '06, (Berkeley, CA, USA), pp. 367–380, USENIX Association, 2006.
[38] Colorado State University, "BGPmon Next generation BGP Monitor." http://bgpmon.netsec.colostate.edu/.

[39] Cooperative Association for Internet Data Analysis, "Detection and analysis of large-scale Internet infrastructure outages." Research Project. http://www.caida.org/funding/dals-satc/.
[40] Cooperative Association for Internet Data Analysis, "Archipelago measurement infrastructure." http://www.caida.org/projects/ark/.
[41] Cooperative Association for Internet Data Analysis, "Macroscopic Topology Measurements." Research Project. http://www.caida.org/projects/macroscopic/.
[42] "Ruby Language." http://www.ruby-lang.org/.
[43] M. Luckie, "Scamper: a Scalable and Extensible Packet Prober for Active Measurement of the Internet," in *ACM SIGCOMM Internet Measurement Conference (IMC)*, 2010.
[44] B. Augustin, T. Friedman, and R. Teixeira, "Measuring load-balanced paths in the Internet," in *ACM SIGCOMM Internet measurement Conference (IMC)*, Oct. 2007.
[45] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. F. Wu, and L. Zhang, "An Analysis of BGP Multiple Origin AS (MOAS) Conflicts," in *Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement*, IMW '01, (New York, NY, USA), pp. 31–35, ACM, 2001.
[46] M. Networks, "List of routing registries," 2012. http://www.irr.net/docs/list.html.
[47] "PeeringDB." https://www.peeringdb.com/.
[48] M. Luckie, B. Huffaker, A. Dhamdhere, V. Giotsas, and K. Claffy, "AS Relationships, Customer Cones, and Validation," in *Proceedings of the 2013 Conference on Internet Measurement Conference*, IMC '13, (New York, NY, USA), pp. 243–256, ACM, 2013.
[49] L. Gao and F. Wang, "Inferring and characterizing Internet routing policies," in *ACM SIGCOMM Internet measurement workshop*, april 2003.
[50] K. Chen, D. Choffnes, R. Potharaju, Y. Chen, F. Bustamante, D. Pei, and Y. Zhao, "Where the sidewalk ends: Extending the Internet AS graph using traceroutes from P2P users," in *ACM CoNEXT*, 2009.
[51] "CAIDA's Macroscopic Internet Topology Data Kit (ITDK)." http://www.caida.org/data/active/internet-topology-data-kit/.
[52] P. Gill, M. Schapira, and S. Goldberg, "Modeling on Quicksand: Dealing with the scarcity of ground truth in interdomain routing data," *ACM Computer Communications Review (CCR)*, 2012.
[53] L. Gao and J. Rexford, "Stable Internet routing without global coordination," *Networking, IEEE/ACM Transactions on*, vol. 9, no. 6, pp. 681–692, 2001.
[54] "American Registry for Internet Numbers (ARIN)." http://www.arin.net.
[55] RIPE Network Coordination Center. http://www.ripe.net.
[56] Digital Envoy, "NetAcuity." http://www.digital-element.net/ip_intelligence/ip_intelligence.html.
[57] B. Huffaker, M. Fomenkov, and K. Claffy, "Automating Inference of Router Locations." Under submission.
[58] A. Robachevsky, "Observations from the routing resiliency measurements workshop," *Internet Society*, 2013.
[59] P. Gill, M. Schapira, and S. Goldberg, "A Survey of Interdomain Routing Policies," *ACM Computer Communications Review (CCR)*, 2014.
[60] Phillipa Gill, "A Survey of Interdomain Routing Policies," 2012. https://www.nanog.org/meetings/abstract?id=1996.
[61] R. Lychev, S. Goldberg, and M. Schapira, "BGP Security in Partial Deployment: Is the Juice Worth the Squeeze?," *SIGCOMM Comput. Commun. Rev.*, vol. 43, pp. 171–182, Aug. 2013.