Does Scale, Size, and Locality Matter? Evaluation of Collaborative BGP Security Mechanisms

Rahul Hiran
Linköping University, Sweden

Niklas Carlsson
Linköping University, Sweden

Nahid Shahmehri
Linköping University, Sweden

Abstract—The Border Gateway Protocol (BGP) was not designed with security in mind and is vulnerable to many attacks, including prefix/subprefix hijacks, interception attacks, and imposture attacks. Despite many protocols having been proposed to detect or prevent such attacks, no solution has been widely deployed. Yet, the effectiveness of most proposals relies on large-scale adoption and cooperation between many large Autonomous Systems (AS). In this paper we use measurement data to evaluate some promising, previously proposed techniques in cases where they are implemented by different subsets of ASes, and answer questions regarding which ASes need to collaborate, the importance of the locality and size of the participating ASes, and how many ASes are needed to achieve good efficiency when different subsets of ASes collaborate. For our evaluation we use topologies and routing information derived from real measurement data. We consider collaborative detection and prevention techniques that use (i) prefix origin information, (ii) route path updates, or (iii) passively collected round-trip time (RTT) information. Our results and answers to the above questions help determine the effectiveness of potential incremental rollouts, incentivized or required by regional legislation, for example. While there are differences between the techniques and two of the three classes see the biggest benefits when detection/prevention is performed close to the source of an attack, the results show that significant gains can be achieved even with only regional collaboration.

I. INTRODUCTION

The Internet is highly susceptible to routing attacks [4], [12]. In almost all types of routing attacks, the attackers rely on vulnerabilities in the Border Gateway Protocol (BGP) to attract traffic that was not intended for them. Often this is achieved through a *prefix attack* or *subprefix attack*, in which the attacker announces itself as the origin of a prefix with the intention of attracting some of the traffic intended for IP addresses belonging to this prefix. Sub-prefix attacks are particularly dangerous as the longest-prefix routing rules implemented on routers always route to the most specific (sub)prefix.

An attack's severity and the complexity of detecting the attack is, to a large extent, determined by the attacker's actions when receiving the hijacked traffic. For example, *black-holing attacks* in which the traffic terminates at the attacker network are relatively easy to detect, as the traffic source may not obtain expected end-to-end responses. In contrast, *imposture attacks*, in which the attacker also impersonates the destination, or *interception attacks*, in which the attacker re-routes the traffic to the destination, are much more difficult to detect.

Unfortunately, despite an increasing number of observed routing attack occurrences [1], [6], [12], [13], it has proven difficult to incentivize operators to invest in existing solutions [4],

TABLE I. EXAMPLES OF SYSTEMS, THE INFORMATION THEY SHARE/USE, AND THE ATTACKS THEY CAN HELP DETECT/PREVENT.

Information shared	Prefix hijack	Subprefix hijack	Inter-ception	Impos-ture	Example solutions
Prefix origin	✓	✓	✗	✗	Route filtering [3], [4], RPKI [21], ROVER [8]
Route path updates	✓	✓	✗	✗	PHAS [20], PrefiSec [15], PG-BGP [16]
Passive measurements	✗	✗	✓	✓	CrowdSec [14]
Active measurements	✗	✗	✓	✓	Zheng et al. [30], PrefiSec [15]

and there is currently no universally deployed solution that prevents hijacking of Internet traffic by third parties [12]. For example, the deployment of crypto-based efforts [18], [21], [27] has been hampered by high deployment costs for network operators [4], [12]. Instead, monitoring of path announcements and the data paths taken by data packets are typically used to identify potential hijacks and other suspicious data paths [14], [16], [20], [28]. With routing paths being determined by the individual routing decisions of many involved operators and other organizations running their own Autonomous Systems (AS) [1], [16], [28], such techniques benefit greatly from information sharing between ASes.

Different types of information can be helpful in the detection of routing attacks. Table I summarizes some of the most commonly proposed information sources for such systems, as well as some example systems and the types of attacks these systems propose to protect against. In this paper, we focus on the first three types and only briefly discuss the fourth type.

A number of important questions arise when considering cooperative information sharing across ASes and other network entities/organizations for the purpose of detecting or preventing routing attacks. For example, how do the detection/prevention rates of the different techniques scale with the number of participants? What is the impact of the size of each participant, or the information available to the participant? And, what is the impact of the location of the participants sharing the information? The latter question may be particularly important as it may help provide insights into the effectiveness of regional government-issued legislation or regional agreements. For example, the United States (US) government or the European Union (EU) may push to have ASes and organizations under their respective jurisdictions share information in order to protect the common interests of each region.

While some of the papers introducing the above example systems have used data-driven analysis to illustrate the power

of large-scale information sharing between large ASes, little attention has been paid to the effect of the geographic locality of each participant. Although many ASes have points-of-presence in many geographic regions, ASes operated by organizations from the same country or geographic region may be more likely to openly or through legislation, for example, share information with each other. Ongoing geographical and political polarization may further contribute to potential location-based participation and sharing restrictions. Motivated by these observations, in addition to analyzing each of the above three questions on their own, this paper places particular focus on the impact of the locality of the participants. Locality is considered both on its own, and also with regards to size-based inclusion within and across regions, as well as with regards to the scale of the (local or global) information sharing alliance.

The main contribution of this paper is a systematic data-driven evaluation of some promising-previously proposed hijack prevention and routing attack detection techniques. In particular, we consider the above outlined questions in the context of three example techniques that share (i) prefix origin information, (ii) route path updates, or (iii) passively collected round-trip time (RTT) information. For our evaluation, we develop a data-driven methodology for each information sharing approach which takes into account the geographic locality (e.g., the region in which the AS is registered) and the relative size (e.g., measured by the number of neighboring ASes) of each of the potential participants. Using real-world topologies and routing information derived from measurement data we then systematically evaluate the impact of each factor, either on its own, or accounting for the geographic locality of the participants, attackers, and victims.

Our results provide insights into the tradeoffs between global and local deployment. While the results highlight the value of detection and prevention close to the source of an attack, we also find cases where regional collaboration may achieve many of the benefits achievable through global deployment. Other interesting findings include the observation that the largest ASes are not always the best at hijack detection when the attacks are from other regions. Instead, collaboration with mid-sized ASes may be beneficial. This is in contrast to the deployment of hijack prevention mechanisms, which benefit significantly from large ASes participating, regardless of whether the deployment is global or regional. Our scale- and size-based evaluation also provides insights into other deployment related issues, including the relative deployment benefits during different phases of an incremental rollout.

Paper outline: Section II provides background and sets the context. The following three sections present our evaluation results for three general classes of collaborative prevention and detection systems. In Section III we evaluate (sub)prefix attack prevention techniques that use prefix origin information, in Section IV we evaluate hijack detection mechanisms that use path announcements, and in Section V we evaluate interception and imposture detection techniques that use passively collected RTT measurements. Finally, related work and conclusions are presented in Sections VI and VII, respectively.

II. BACKGROUND

BGP works well in normal circumstances. However, inherent vulnerabilities with the protocol enable routing attacks.
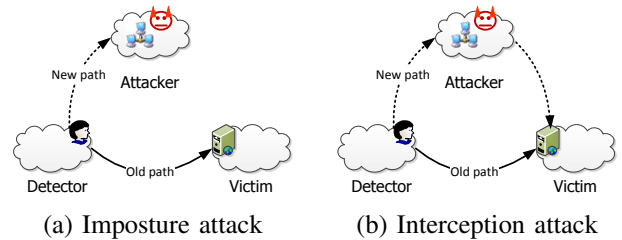

(a) Imposture attack   (b) Interception attack

Fig. 1.   Imposture and interception examples.

A. Routing attacks

A major vulnerability in BGP is its inability to validate the allocation of prefixes to ASes. This makes it difficult to detect when an attacker AS announces one or more prefixes allocated to other network(s). In a *prefix hijack* the attacker announces a prefix (e.g., a.b.c.d/16) that is actually allocated to a different AS. Depending on ASes' relationships and how the AS-PATH is propagated between ASes such attacks may attract (hijack) more or less traffic. In a *subprefix hijack* the attacker announces a subprefix (e.g., a.b.c.d/24) of a larger prefix (e.g., a.b.c.d/16). Due to the longest-prefix matching rule used by routers, these attacks may be particularly effective in hijacking traffic.

All the above types of attack may lead to one of several outcomes. For example, in a *blackholing attack* the attacker simply drops the traffic that it attracts. Figure 1 illustrates two more difficult attacks to detect. In an *imposture attack* (Figure 1(a)), the attacker impersonates the intended destination for the traffic and in an *interception attack* (Figure 1(b)) the attacker redirects the traffic to its intended destination, possibly after making a copy or modifying the data, for example. These attacks are particularly stealthy when the users originating the traffic receive uninterrupted service.

B. Collaborative information sharing

Various systems have been proposed to detect, mitigate, or prevent routing attacks and other unwanted routing incidents. These systems typically rely on collaborating ASes sharing different information. For our analysis we focus on three broad classes of techniques that share and/or use the first three types of information in Table I. They correspond to prefix origin information, route path updates, and passively collected RTTs. Route path updates can easily be collected at individual routers or at the AS level, and then shared with other ASes. RTTs can easily be passively collected and shared by almost any network entity [14]. In the following we describe how the different systems that we evaluate here use the shared information.

C. Hijack prevention using prefix origin

Ideally, a hijack prevention mechanism should prevent an AS from accepting and propagating bogus route announcements. If implemented widely, such mechanisms could then prune away bogus route announcements close to the source and prevent an attacker from reaching more ASes and users.

Previously proposed mechanisms that can provide hijack prevention include prefix filtering [4], PG-BGP [16], RPKI [8], and ROVER [17]. Many of these techniques build a trusted and formally verifiable database of prefix-to-AS pairings between the IP prefixes and the ASes that are allowed to originate them. If used correctly, routers implementing RPKI [8] and

ROVER [17] can be assumed to pick the right origin and avoid propagating the wrong origin for prefixes. To achieve this, two conditions must be satisfied. First, the AS that wants to defend itself against route hijacks on its prefixes, or its provider AS, must register a prefix-to-AS mapping with RPKI or ROVER. Second, as part of their validation of prefix-to-AS mappings, relying ASes must successfully retrieve and check these AS-to-prefix mappings from the RPKI or ROVER records.

With PG-BGP [16], the acceptance of suspicious routes is delayed, and routes are accepted and propagated only after a certain threshold time duration has passed. Since suspicious routes are typically short-lived [25], the performance of PG-BGP is usually similar to that of RPKI and ROVER. Therefore, for the purpose of our evaluation, we only simulate the performance of RPKI and ROVER.

D. Control-plane based anomaly detection

There are several works that are based on control-plane data for the detection of anomalies in BGP routing, including PHAS [20], PrefiSec [15], and PG-BGP [16]. While PHAS and PG-BGP aggregate all information centrally, PrefiSec distributes computing and detection across participants. Otherwise, the approaches are relatively similar. For each prefix, these protocols track the origin ASes observed by its participants and raises alerts when there are changes. The common idea leveraged by all these protocols is that an IP prefix should be originated by a single AS. An IP prefix originated by more than one AS results in a Multiple Origin AS (MOAS) conflict. While some MOAS are legitimate and can be observed over long time periods [29], a newly-detected MOAS conflict can be an indication of a potential prefix hijack. By keeping track of the AS-to-prefix mappings observed in AS-PATH announcements, these protocols can flag new potential MOAS cases. Naturally, as more ASes participate and share their observed path announcements, the system will have more complete AS-to-prefix mappings.

E. Route anomaly detection using passive measurements

The examples in Figure 1 illustrate why imposture and interception attacks may be particularly difficult to detect without observing the actual data path or the impact these changes have on the RTTs. Both active traceroute-based anomaly detection techniques [15], [30] and passive RTT-based anomaly detection techniques [14] have been proposed. While we will focus on the use of passive measurements, we note that the approaches in general are fairly similar. For example, Zheng et al. [30] use changes in the number of hops in the traceroute paths to identify potential hijacks, while Hiran et al. [14] use changes in the RTTs to identify potential anomalies.

In both types of systems measurement information from multiple sources is shared to provide stronger evidence and more accurate flagging of suspicious events. For example, in CrowdSec [14] clients or other network entities (e.g., middleboxes) share RTT outlier information and collaboratively identify prefixes with many affected clients, so as to identify potential routing anomalies. For collaborative detection the system uses statistical tests based on binomial hypothesis testing. One of the main advantages of using passive measurements is that, in contrast to active measurements such as traceroutes, they do not add additional traffic overhead.

III. EVALUATING HIJACK PREVENTION TECHNIQUES

Several studies have suggested that there are significant benefits to deploying hijack detection and prevention mechanisms on several large ASes across the world. However, global deployment that spans multiple geographic regions and jurisdictions is non-trivial and may not be practical due to political and economic reasons. It may be more practical to push or incentivize the deployment within a geographic region such as the US or EU. For example, governmental legislation or other regional mechanisms may be used to push or incentivize agreements between ASes within a region.

In this paper, we evaluate and compare the benefits and drawbacks of deploying three different general classes of prevention and detection techniques regionally versus globally. For each class of techniques we simulate the effectiveness of the general technique when different subsets of potential candidate participants employ the technique and share information between each other. Within this context, we then answer questions related to the impact of locality and size of the participants, as well as the number of participants. For example, what is the impact of the number of ASes that deploy the hijack prevention mechanisms, either from a specific region (e.g., North America or Europe) or globally? And, what is the impact of size of the ASes that deploy the hijack prevention mechanisms from a specific region or globally?

In this section we answer the above questions in the context of hijack prevention mechanisms such as route filtering [3], [4], RPKI [21], and ROVER [8].

A. Simulation-based Evaluation Methodology

For our simulation-based evaluation, we modified and extended the existing BSIM [16] simulator. BSIM simulates route propagation using the standard Gao-Rexford model [7], which captures the behavior of the economy-driven policies used in practice [11]. The model distinguishes between *customer-provider* relationships (where the customer AS pays its provider) and *peer-peer* relationships (where two ASes often agree to transit each other's traffic for free). In particular, the model assumes that ASes use a routing policy in which customer routes may be exported to all neighboring ASes, but routes learned from peers or providers are exported only to the customers. In addition, the policy prefers customer routes over peer routes (since they bring revenue) and peer routes over provider routes (since provider routes cost money). In cases of multiple tied routes, the routes with the shortest AS paths are chosen. Finally, for the purpose of the simulations, if there are still ties, these ties are broken (arbitrarily) by picking the route over the AS with the lowest AS number.

We extracted the Internet AS-level topology and AS relationship information for every pair of neighboring ASes from public data [5]. We use a snapshot from August 2015, which contains 51,507 ASes and 199,540 AS relationships.

For evaluating hijack prevention mechanisms, we simulate how the routes would propagate in the presence of hijack prevention mechanisms compared to the case when these mechanisms are not present. We measure the fraction of ASes that end up forwarding packets along the correct path in both scenarios and report the percentage increase in the number of

ASes that choose the correct origin. To calculate the percentage increase we first simulate each example attack when no ASes deploy the prevention mechanism and when a random subset of ASes deploy the mechanism, respectively, and then report the average increase in the number of ASes that route to the correct destinations when the prevention mechanism was deployed. In each example scenario, we perform simulations by randomly choosing victim and attacker ASes from selected example regions. Across all scenarios, we randomly picked N ASes to deploy the mechanism from the set of ASes with at least X neighboring ASes, and reported the average over 500 simulations per scenario (with 95% confidence intervals for the average). The use of threshold is in part motivated by larger ASes (with many neighbors) being more likely to have the resources to deploy hijack prevention mechanisms [10]. Here, the degree threshold X is used to bias the size of the individual participants and the parameter N captures the size (scale) of the alliance as a whole.

To compare different deployments, we use locality-, size-, and scale-based criteria to randomly pick subsets of the nodes on which to implement the mechanism. In all simulations, victim nodes are selected at random and the reported metrics are calculated over all nodes in the network.

As is common practice, for our evaluation we varied one parameter at a time, while keeping all the other parameters constant. Our default degree threshold $X = 20$ was selected to map to an intermediate value in the range of interest (0-50), and the default alliance size N was selected to be equal to the number of ASes with at least 50 neighbors.

Before presenting our results, it should be noted that the simulations have limitations. First, the AS relationship data used for the simulations is not perfect and does not take into account more complex AS-to-AS relationships. For example, two ISPs may interconnect at multiple peering points and have different types of relationships at each point [5]. Second, not all network operators follow the standard rules for route export. However, it is believed that there are few exceptions [11].

B. Global Baseline: Scale and Size

For reference, we first present results when the participating ASes are selected from the global set of ASes. Figure 2 summarizes these results. Figure 2(a) shows the percentage improvement in the number of ASes that chose the correct origin, as a function of the number of participating ASes. With our default threshold $X = 20$, the right-most points correspond to the case when all 2,626 ASes with at least 20 neighbors participate. In comparison, Figure 3 shows the same plot for ASes in North America (NA), the European Union (EU), and the rest of the world (all ASes excluding those in NA and EU).

Referring to the global deployment results (Figure 2(a)), all regions observe significant advantages from higher participation. For example, with 500 random participants we observe an average improvement of more than 15% across all victim-attacker pair scenarios. In comparison, hen all ASes with degree of at least 20 participate the improvements are consistently above 45%. While overall numbers are lower when only local ASes participate (Figure 3), we note that local deployment is important when protecting against attacks from within the region. This is demonstrated by the higher percentage of improvements when the attacker is in the region deploying the security mechanism.

Figure 2(b) shows the percentage improvement as a function of the threshold degree X. With our default alliance size $N = 1,093$, the righ-most points correspond to the case when all ASes with a degree of at least 50 participate.

From these results it is clear that the high-degree ASes are the ones that offer the most protection. For example, if all the 1,093 top-ASes with more than 50 neighbors participate we observe improvements of more than 40% for all victim-attacker scenarios. This shows the importance of getting the large ASes onboard in these deployment efforts. The general observation that collaboration by a few large ASes can provide much of the protection is not new. Similar observations have been made by Gersch et al. [9] and Karlin et al. [16], for example. In this work, we take this analysis one step further and consider the impact of regional deployment.

Figure 4 presents location-based results for when only ASes in a certain region deploy the prevention mechanism, and where ASes deploying the mechanism are selected based on their degree. We note that regional deployment can provide similar improvements as in a global deployment when the attacker is local. The improvements are noticeably lower when the attackers are located in other regions. For example, the percentage gain in ASes choosing the correct origin for attackers in NA is greater when ASes in NA deploy the prevention mechanisms compared to the gains when ASes in other regions deploy the mechanisms. These results illustrate that enforced deployment of these mechanisms may be a good way for regions to clean up their own networks.

The locations of victim networks play a smaller role. Although all networks would benefit from such a deployment, the local networks would not gain much more protection than external networks. These mechanisms should perhaps best be seen as mechanisms for the greater good, with the results showing that there is great incentive for governments and network operators to come together to help ensure that prevention mechanisms are deployed on a large scale.

C. Location-based Discussion: Key Findings

It is often stated that you should *keep your friends close and your enemies closer*. Our results highlight that this is also an important lesson in today's networks. First, starting with our friends, our results show that there are substantial gains from local deployment, regardless of where the attacks come from. For example, if all ASes with a degree of at least 20 deploy these mechanisms we observe a 30% gain for the NA-based victims regardless of attacker region (e.g., Figure 3(a)). In EU (Figure 3(b)) the corresponding gain is 20%.

Second, considering the attackers, our regional results clearly show that detectors close to the attackers help the most. For example, when ASes in NA deploy hijack prevention schemes, the damage from hijack attacks originating from ASes in this region can be significantly controlled in all regions. These results show that the largest benefits come with global deployment, and that there may be benefits to subsidizing or otherwise incentivizing international partners to implement these mechanisms too.

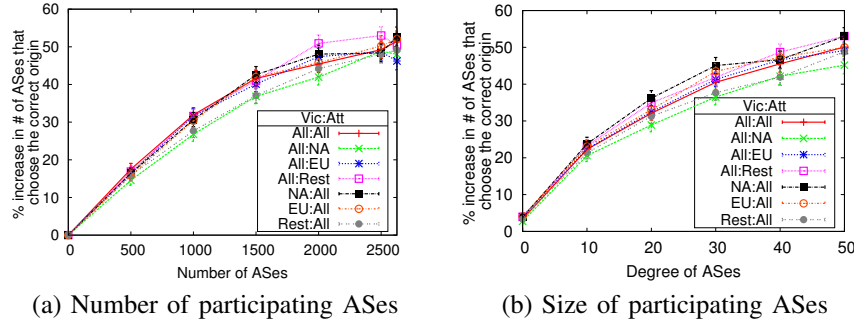(a) Number of participating ASes   (b) Size of participating ASes

Fig. 2. The average percentage improvement in the number of ASes that choose the correct origin when different subsets of the global set of ASes participate.
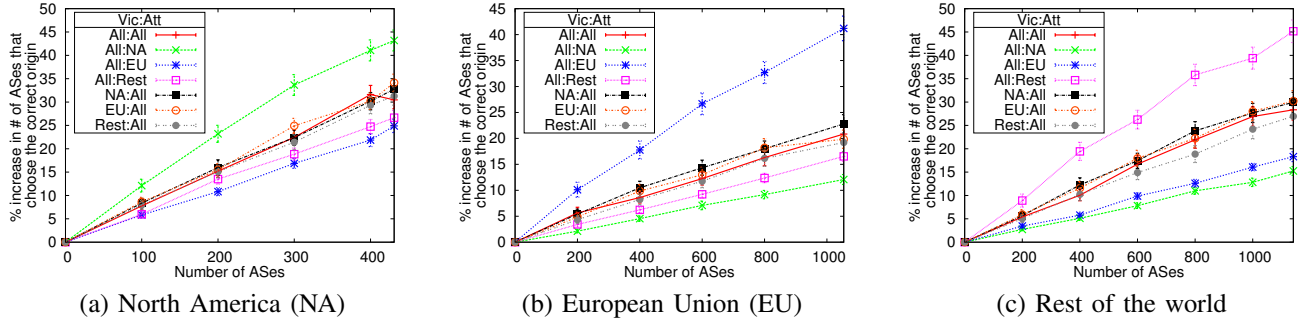

(a) North America (NA)   (b) European Union (EU)   (c) Rest of the world

Fig. 3. Impact of number of participating ASes, when ASes are selected from a particular geographical region or the "rest of the world".
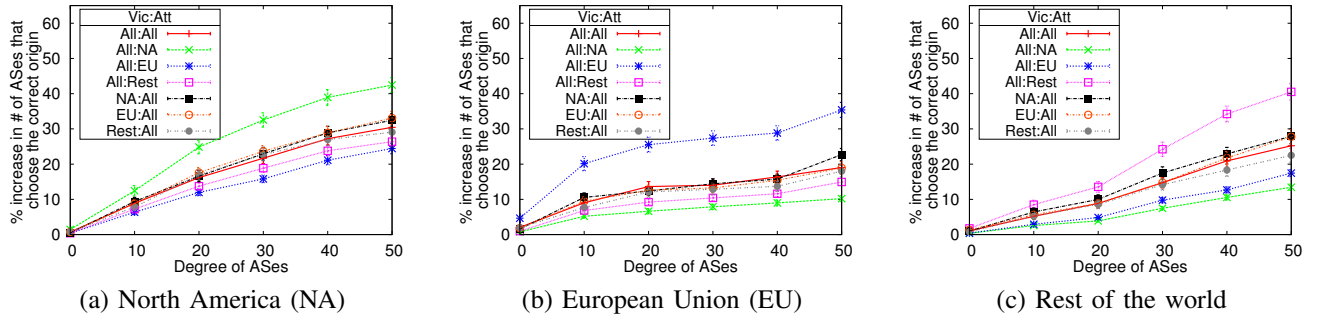

(a) North America (NA)   (b) European Union (EU)   (c) Rest of the world

Fig. 4. Impact of the degree threshold of the participating ASes, when all are selected from a geographic region or the "rest of the world". For these figures, we choose $N = 207$, $N = 571$, and $N = 315$, respectively.

Overall, the percentage gain when hijack prevention mechanisms are deployed by all (roughly 2,500) ASes around the world with a degree of at least 20 varies between 40% and 50% for different combinations of victim-attacker regions (Figure 2(a)). If only 500 random ASes around the world deploy the mechanism, the gain is roughly 15%. By comparison, when the 431 ASes with a degree of at least 20 in NA deploy the hijack prevention mechanisms, the percentage gain varies between 23% to 43%, depending on which victim-attacker pair combination is considered (Figure 3(a)). The higher numbers partially reflect the big impact of the NA-based ASes, many of which are high-degree ASes with peering points around the globe, but also demonstrate the value of regional deployment to help protect against hijack attacks.

IV. EVALUATING HIJACK DETECTION MECHANISMS

A. Methodology and Datasets

To evaluate hijack detection mechanisms based on AS-PATH updates we have extended and modified a framework that was previously used to evaluate alert rates for PrefiSec [15] to account for ASes and their locality. While the evaluation framework is designed for PrefiSec, the results presented here also apply to PHAS [20] and PG-BGP [16]. Given the same

information, these systems' detection rates are the same. The main differences between these systems are their communication overhead and where the processing is performed.

For our analysis, we collected the RIB files and AS-PATH announcements observed at all six routeviews servers active during the time of the China Telecom incident [13], on April 8, 2010, when China Telecom announced origin for 50,000 prefixes originated by other ASes. Using announcements from around the time, we compare differences and similarities of the detection rates during an actual attack.

Focusing on a two-week window around the time of the incident, we first used the RIB files from April 1 and a warm-up period to initialize the AS-to-prefix mappings seen by different selected subsets of ASes. Of particular interest here is the degree (size) and locality of the collaborating ASes. In our evaluation, we consider sets of collaborating ASes selected from NA, EU, the "rest of the world" (reference point, rather than a region), and from the global set. In contrast to the original evaluation frameworks, which treated the routeviews servers as the participants [15], we use the AS information of the ASes contributing announcements to the routeviews servers and AS-to-region mappings to identify subsets of information seen by different subsets of collaborating ASes.

In total, the six routeviews servers have 100 vantage points that belong to 73 unique ASes. Of these, 38 are NA-based, 21 EU-based, and 14 map to other geographic regions.

For each of the subsets of collaborating ASes that we chose, we then look at each day in the time window and simulate and report the number of prefixes and origins, respectively, that the ASes in the subset would not have seen prior to that day. These two metrics directly measure the number of cases that must be flagged (and further investigated) as potential prefix and sub-prefix attacks, respectively.

B. Global Baseline

As a baseline, we first present results for when the collaborating ASes are selected globally. Figure 5(a) shows the number of alerts raised for both "new prefixes" (possible subprefix hijacks) and "new prefix origins" (possible prefix hijacks) announced during the incident (on April 8) as a function of number of collaborating ASes. We also include separate lines for the number of alerts of these two types raised due to announcements made by China Telecom.

We see that the number of alerts for possible prefix hijacks increases with the number of collaborating ASes, and that 40,575 alerts (for both prefix and subprefix hijacks) are raised during the day of the attack if all the nodes collaborate. With the exception of a few "new prefixes" and "new prefix origins", almost all alerts are due to the China Telecom announcements associated with the incident, which caused traffic for these prefixes and subprefixes to be hijacked.

Only a few ASes are needed to detect the majority of the subprefix hijacks ("new prefixes"). This result can be explained by subprefixes being propagated to almost all ASes due to more specific prefixes being preferred. For prefix attacks ("new origin") additional ASes are much more beneficial, with some diminishing returns after reaching 40 ASes. This happens because ASes during these instances become divided into two groups: ASes that continue routing to the victim network and ASes that choose to route to the attacker network. Thus, additional collaborating ASes increases the chance that conflicting origins are detected and hijack alerts are raised.

Figure 6 puts the above numbers in perspective, showing the number of alerts for the days before and after the attack. In addition to being orders of magnitude lower than during the day of the incident, the flatter "new origin" curves suggest that the "new origin" announcements during these days propagated somewhat further than the China Telecom announcements.

Figure 5(b) shows the number of alerts as a function of the degree threshold to be included in the alliance. For every threshold, 10 ASes with a degree of at least X are selected at random. Here, the right-most displayed threshold is picked so that the selection set include exactly 10 ASes, and the following points (moving to the left) are picked so as to roughly double the selection set for each point. The degree threshold of 1 is included as a reference point.

The figure shows that the number of alerts for the China Telecom incident is higher when the degree threshold is small, and the number of alerts is quite low when large ASes collaborate. This is a very interesting observation as much prior work has suggested collaboration between the largest ASes,

but it can be partially[1] explained by most of the high degree ASes being NA-based. For example, of the ASes with a degree greater than 1,174, all but one (i.e., 9 of 10) are NA-based, and when the threshold is 646, there are 18 NA-based and 2 EU-based. However, these NA-based ASes do not have as good a vantage point of the China-based incident, with only a subset of the paths propagating to these ASes. With a lower degree threshold more ASes from outside NA and EU will be included, improving the results. This illustrates that the vantage points offered by global collaboration can be more valuable to the prefix hijack detection than having only the large ASes collaborate. Similarly, multi-hop BGP peering can also help. The detection numbers for subprefix attacks ("new prefixes") are less dependent of the AS degree (size) and locality; again, indicating their wider propagation.

C. Location-based Analysis

We now discuss the benefits of regional collaboration for hijack detection. Figure 7 shows the number of alerts as a function of number of ASes for different regions. For all of the three regions (NA, EU, and "rest of the world"), the number of alerts increases as more ASes share information. If all NA-based ASes collaborate there are 22,178 alerts (13,214 "new origin" and 8,964 "new prefix"). Sharing among all EU-based ASes raises 10,829 (3,620+7,209) alerts and sharing among all the ASes in the "rest of the world" category would raise 36,328 (27,280+9,048) alerts. Whereas the sub-prefix detection ("new prefix") is similar for the different regions, the differences in total alerts are substantial. For example, despite there being far fewer ASes in the "rest of the world" category, this category has the highest detection rate. The main reason for this is that many of these ASes have more vantage points closer to China Telecom than NA-based and EU-based ASes may have, and therefore have better visibility of the route announcements made by China Telecom. This observation mirrors the insights provided by our hijack prevention results (Section III) that show that ASes deploying protection mechanisms close to the attacker provide the best protection.

While none of the regional collaborations performs as good as global collaboration, the value of regionally deployed solutions should not be underestimated, especially as there is no solution that has seen widespread deployment yet. These results show that careful regional deployment, possibly with a few complementing ASes from other regions, may provide a significant step in the right direction.

Figures 8(a) and 8(b) show the number of alerts as a function of the degree threshold for regional collaborations in NA and EU, respectively. As for the global results, for each degree threshold, we randomly pick 10 ASes per alliance.

We again observe stronger degree (size) dependence for prefix hijack detection ("new origins") than for subprefix hijack detection ("new prefixes"). While the large ASes in NA in general provide more alerts than the smallest ASes in NA, it is very interesting that the very top ASes see a drop in the number of alerts they raise. It is also interesting that the large ASes in EU detect fewer attacks than the smaller ASes in EU. As the above ASes are in the same region, our previous explanations (in Section IV-B) regarding the relative differences in coverage

---

[1]Additional explanation will be provided in the next subsection.
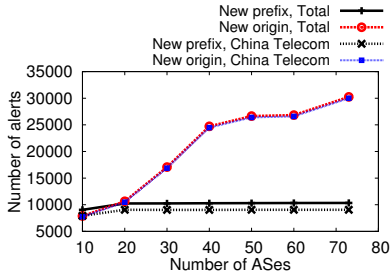
(a) Number of participating ASes
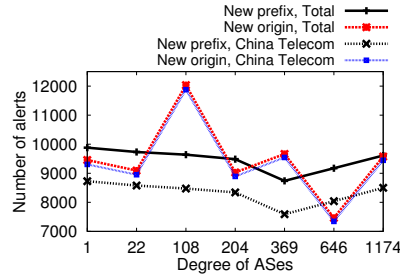
(b) Size of participating ASes

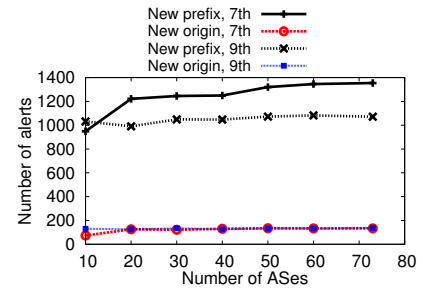Fig. 5. Average number of alerts raised when global ASes collaborate the day of the China Telecom incident.
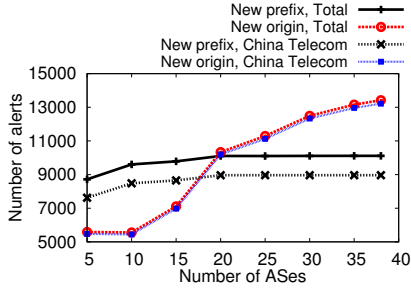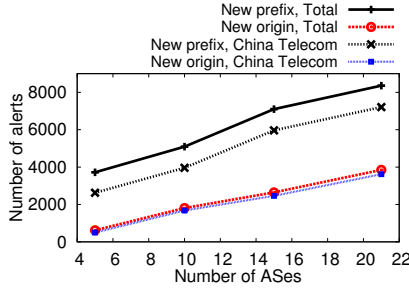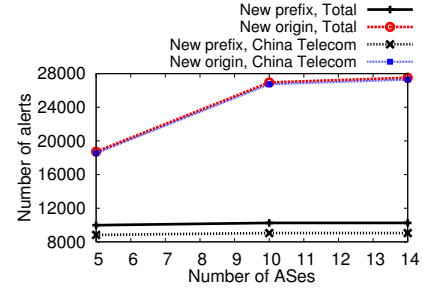

Fig. 6. Average number of alerts raised when global ASes collaborate the day before (April 7) and after (April 9) the incident.

(a) North America (NA)

(b) European Union (EU)

(c) Rest of the world

Fig. 7. Number of alerts during the day of the incident (April 8, 2010) for different sizes of regional collaborations.

seen by ASes in *different regions* no longer apply here. In the *same region*, the size-based differences may instead be related to the standard route export policy. In particular, malicious routes (learnt from a peer or provider) are typically exported only to customers. Therefore, malicious routes learnt by mid-tier ASes may not reach their providers (typically large ASes).

V. INTERCEPTION AND IMPOSTURE DETECTION

A. Methodology and Datasets

To provide insights into the impact of regional collaboration for detecting interception and imposture attacks, we have extended the evaluation of CrowdSec [14] to account for locality of the collaborating network entities. CrowdSec is designed to raise alerts about RTT anomalies and help detect interception and imposture attacks. In the case of an interception attack (Figure 1(b)) the RTTs typically increase during an attack, whereas the RTTs during an imposture attack (Figure 1(a)) can either increase or decrease, depending on the relative locality of the attacker, victim, and detector.

In CrowdSec the end clients passively collect RTT measurements while in contact with different candidate victim IP addresses (or prefixes). The client applies an outlier detection test to raise an alert if the new RTT measurement deviates significantly from previously observed RTT measurements. These alerts are shared with other CrowdSec clients, and the individual alerts are combined using statistical test methods such as a binomial test that takes into account the likelihood of N clients observing significant deviations in RTT measurements to the same prefix, given past observations [14].

For the evaluation presented here, passively collected RTT values are simulated by extracting RTTs from (active) traceroute measurements performed by PlanetLab[2] nodes as part

---
[2]PlanetLab, https://www.planet-lab.org/

of the iPlane [23] project. In particular, we use daily RTT measurements associated with 106 NA-based nodes, 79 EU-based nodes, and 36 nodes located in other parts of the world. For the most part we use a month's worth of training data (e.g., 278,690 successful traceroutes during July 2014) and evaluate the performance of different detection techniques for the following week, during which we simulate different attack combinations. In total, we simulate 15,279 interception attacks and 62,576 imposture attacks per set of sample detectors, and report results averaged over 10 such sample sets. While we only present interception results, the results for imposture attacks are similar. In each simulation, detector nodes and affected nodes are selected randomly within each region.

B. Global and location-based evaluation

Figure 9 shows a comparison of tradeoffs in detection rate during a simulated attack (y-axis) and the false alert rate under normal circumstances (x-axis), when all global vantage points are collaborating (Global) and when only those in North America (NA) or Europe (EU) collaborate. We include results for when all (100%), half (50%), or none (0%) of the potential detector nodes are affected. The case when no nodes (0%) are affected is included only as a reference point, and captures the false positive rates during normal circumstances.

Interestingly, the 106 NA-based nodes achieve a tradeoff that is almost as good as the larger global collaboration (with 221 nodes). For example, a global collaboration that allowed a false alert rate of 10^{-2} would achieve a detection rate of 80%, if 50% of the nodes were affected. In the same scenario, the NA-based nodes achieve a 70% detection rate and the EU-based nodes achieve a 40% detection rate.

Note, however, that the size of the collaboration may play a big role. In Figure 10 we present a regional comparison while keeping the number of detector nodes fixed at 20 and
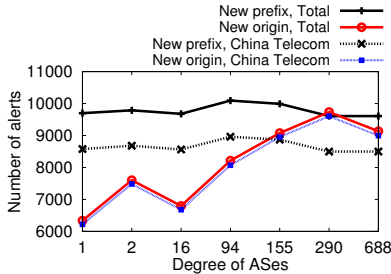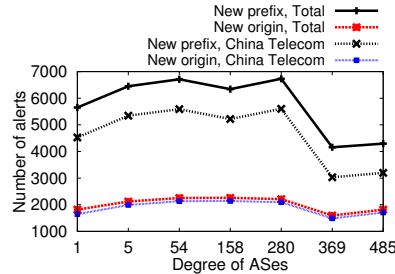
| (a) North America (NA) | (b) European Union (EU) |

Fig. 8. Impact of the size of the participating ASes on the number of alerts. For each degree threshold we choose 10 ASes with a degree equal to or greater than the applied threshold.
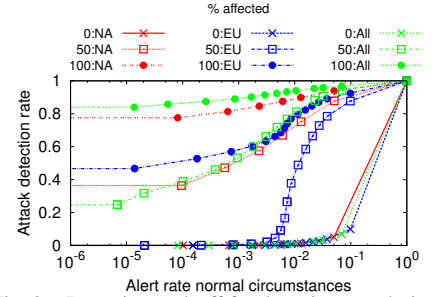
Fig. 9. Detection tradeoff for detection rate during attack and false alert rate during normal circumstances for varying percentages of affected nodes.

30. It turns out that NA-based nodes provide much better detection than EU-based nodes, even when taking alliance size into account, and in fact outperform a global alliance. Part of the reason for these differences may be differences in the variability of RTTs. Another possible contributing factor that we have discovered is that some EU-based routes (even between two EU-based nodes) go through NA even under "normal" circumstances. In such cases, attacks by networks outside EU may not result in noticeable changes in the RTTs.

C. Scale of collaboration

In general, regardless of the locality of the alliance, we have observed significant advantages to larger alliances. This is illustrated in Figure 11. Here, we show the alert-rate tradeoff for collaborations of different sizes when including nodes that are randomly selected from all global nodes (Figures 11(a)) vs. only North America (Figures 11(b)). Related to scale, it should also be noted that there are benefits to larger numbers of RTT measurements, as this helps to filter out anomalies. While region-based analysis of this aspect is omitted here, we refer the interested reader to our global results [14].

VI. RELATED WORK

A large number of security mechanisms have been proposed to secure Internet routing. As described in Section II, this includes prefix hijack prevention mechanisms based on prefix filtering [3], [4], crypto-based solutions such as RPKI [21] and ROVER [8], hijack detection mechanisms based on changes in prefix origins observed in AS-PATH announcements (e.g., PHAS [20], PrefiSec [15], and PG-BGP [16]), and route hijack detection mechanisms using either passive RTT measurements [14] or active traceroute measurements [15], [30]. Rather than proposing new mechanisms, we evaluate the effectiveness of three broad classes of such mechanisms when they are only partially deployed. We place particular focus on the geographic locality of the collaborating ASes or network entities, while also considering the impact of the collaboration scale and the size of participating ASes.

While partial deployment of BGP security mechanisms has been considered in prior literature [2], [10], [16], [22], the geographic location of participants is almost always ignored. Instead, carefully selected ASes have typically been used to demonstrate the potential of the individual techniques. For example, Avramopoulos et al. [2] demonstrate good protection of a participant's outgoing and incoming traffic using only the top-5 tier-1 ASes in the world. Others have relied on the top-tier ASes to demonstrate the effectiveness of PG-BGP [16],

path validation protocols such as S-BGP and BGPSec [22], and incentive strategies for deployment of S*BGP [10]. None of these works consider the impact of locality of the ASes that are deploying the security mechanisms.

We are not the first to study the impact of the number of participating ASes [9] or their node degree [24]. For example, Suchara et al. [24] analyze security gains as a function of increasing the node degree of the ASes that use a BGP security mechanism that filters malicious routes. Similarly to our results, they find significant benefits to deploying the mechanism at high-degree ASes at the core of the Internet. Gersch et al. [9] analyze the effect of increasing the number of ASes using attack prevention techniques. Their results nicely show how the average number of polluted ASes decreases as the number of participating ASes (with higher degrees) increases. Again, none of these works consider which geographic region each AS maps to. This can be an important factor when it comes to legislation and other political incentives.

Much work has also been done to understand the slow adoption of RPKI and other solutions [10], [26]. Other orthogonal but interesting work in this domain has designed AS reputation systems that use control-plane information to capture short-lived routes often used by malicious ASes [19].

Finally, the original simulation framework used in Section III has also been used by Karlin et al. [16]. For this part, we extend the simulator to take into account the geographic locations of the attackers, victims, and collaborating participants. In Sections IV and V we extend and generalize our prior evaluation frameworks for PrefiSec [15] and CrowdSec [14]. Again, neither of these systematically evaluates the value of scale and size in the context of locality-restricted collaboration. This paper evaluates three such broad classes of mechanisms.

VII. CONCLUSIONS

Despite BGP's vulnerabilities and increasingly many routing attacks, no universally deployed security solution to such attacks exists. Using simulations based on real measurement data we have presented a systematic evaluation of three broad classes of prevention and detection techniques. We have focused on the impact that regional, rather than global, deployment could have on their ability to prevent/detect attacks, as well as the impact of AS size of the (regional or global) participants and the number of ASes that deploy the techniques. While prefix hijack prevention (Section III) and detection (Section IV) benefit greatly from deployment close to the source of an attack, it is encouraging to see cases with
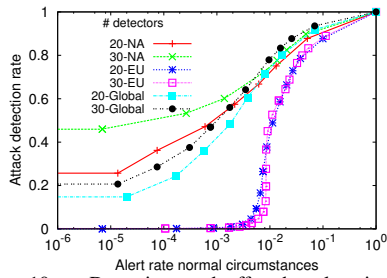
Fig. 10. Detection tradeoff, when keeping the number of detectors fixed. Here, 50% of the nodes are assumed to be affected.
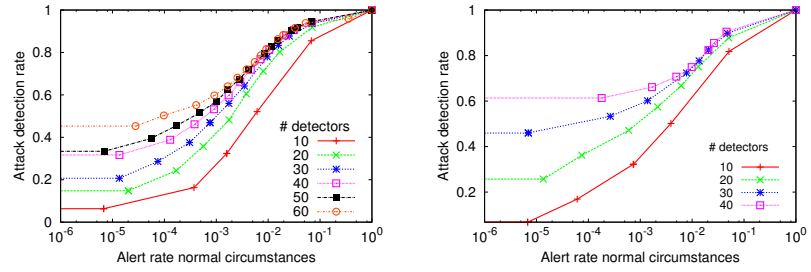

(a) Global (b) North America (NA)

Fig. 11. Detection tradeoff for detection rate during attack and false alert rate during normal circumstances for varying numbers of detectors. Here, 50% of the nodes are assumed to be affected.

all three classes of techniques where regional deployment provides substantial benefits. We even find some cases where regional deployment achieves most of the benefits achievable through global deployment, and note that regional deployment with carefully selected participants (e.g., based on AS size) can outperform global deployments that are less carefully planned. Another interesting observation is that the largest ASes can provide worse detection than mid-sized ASes, which may see a richer set of bogus announcements. This contrasts to deploying hijack prevention mechanisms, for which large ASes appear to provide the greatest benefit. The best AS selection may therefore depend on if the system is designed for prevention or detection. We have focused on one class of techniques at a time. Interesting future work could weigh the benefits of the different approaches against each other for different collaboration constellations.

ACKNOWLEDGEMENTS

REFERENCES

[1] ARNBAK, A., AND GOLDBERG, S. Loopholes for circumventing the constitution: Unrestrained bulk surveillance on Americans by collecting network traffic abroad. In *Proc. HOTPETS* (Jul. 2014).

[2] AVRAMOPOULOS, I., SUCHARA, M., AND REXFORD, J. How small groups can secure interdomain routing. Tech. rep., Princeton University, Nov. 2007.

[3] BATES, T., GERICH, E., JONCHERAY, L., JOUANIGOT, J.-M., KARRENBERG, D., TERPSTRA, M., AND YU, J. Representation of IP routing policies in a routing registry. RFC 1786 (Informational), Mar. 1995.

[4] BUTLER, K., FARLEY, T., MCDANIEL, P., AND REXFORD, J. A survey of BGP security issues and solutions. *Proc. IEEE 98*, 1 (Jan. 2010), 100–122.

[5] DIMITROPOULOS, X., KRIOUKOV, D., FOMENKOV, M., HUFFAKER, B., HYUN, Y., CLAFFY, K., AND RILEY, G. As relationships: Inference and validation. *SIGCOMM CCR 37*, 1 (Jan. 2007), 29–40.

[6] DYN RESEARCH. Pakistan hijacks YouTube, 2008.

[7] GAO, L., AND REXFORD, J. Stable Internet routing without global coordination. *ACM SIGMETRICS 28*, 1 (Jun. 2000), 307–317.

[8] GERSCH, J., AND MASSEY, D. ROVER: route origin verification using DNS. In *Proc. IEEE ICCCN* (Jul/Aug. 2013).

[9] GERSCH, J., MASSEY, D., AND PAPADOPOULOS, C. Incremental deployment strategies for effective detection and prevention of BGP origin hijacks. In *Proc. IEEE ICDCS* (Jun. 2014).

[10] GILL, P., SCHAPIRA, M., AND GOLDBERG, S. Let the market drive deployment: A strategy for transitioning to BGP security. In *Proc. ACM SIGCOMM* (Aug. 2011).

[11] GILL, P., SCHAPIRA, M., AND GOLDBERG, S. A survey of interdomain routing policies. *SIGCOMM CCR 44*, 1 (Jan. 2014), 28–34.

[12] GOLDBERG, S. Why is it taking so long to secure Internet routing? *ACM Queue 12*, 8 (Oct. 2014), 327–338.

[13] HIRAN, R., CARLSSON, N., AND GILL, P. Characterizing large-scale routing anomalies: A case study of the China telecom incident. In *Proc. PAM* (Mar. 2013).

[14] HIRAN, R., CARLSSON, N., AND SHAHMEHRI, N. Crowd-based detection of routing anomalies on the Internet. In *Proc. IEEE CNS* (Sep. 2014).

[15] HIRAN, R., CARLSSON, N., AND SHAHMEHRI, N. PrefiSec: A distributed alliance framework for collaborative BGP monitoring and prefix-based security. In *Proc. ACM CCS WISCS* (Nov. 2014).

[16] KARLIN, J., FORREST, S., AND REXFORD, J. Pretty good BGP: Improving BGP by cautiously adopting routes. In *Proc. IEEE ICNP* (Nov. 2006).

[17] KENT, S. An infrastructure supporting secure Internet routing. In *Public Key Infrastructure*, vol. 4043. 2006, pp. 116–129.

[18] KENT, S., LYNN, C., AND SEO, K. Secure Border Gateway Protocol (S-BGP). *IEEE Journal on Selected Areas in Communications 18*, 4 (Apr. 2000), 582–592.

[19] KONTE, M., PERDISCI, R., AND FEAMSTER, N. ASwatch: An AS reputation system to expose bulletproof hosting ASes. In *ACM SIGCOMM* (Aug. 2015).

[20] LAD, M., , PEI, D., WU, Y., ZHANG, B., AND ZHANG, L. PHAS: A prefix hijack alert system. In *Proc. USENIX Security* (Jul/Aug. 2006).

[21] LEPINSKI, M., AND KENT, S. An infrastructure to support secure Internet routing. RFC 6480 (Informational), Feb. 2012.

[22] LYCHEV, R., GOLDBERG, S., AND SCHAPIRA, M. BGP security in partial deployment: Is the juice worth the squeeze? In *Proc. ACM SIGCOMM* (Aug. 2013).

[23] MADHYASTHA, H., ISDAL, T., PIATEK, M., DIXON, C., ANDERSON, T., KRISHNAMURTHY, A., AND VENKATARAMANI, A. iPlane: An information plane for distributed services. In *Proc. OSDI* (Nov. 2006).

[24] SUCHARA, M., AVRAMOPOULOS, I., AND REXFORD, J. Securing BGP incrementally. In *Proc. ACM CoNEXT* (Dec. 2007).

[25] VERVIER, P.-A., THONNARD, O., AND DACIER, M. Mind your blocks: On the stealthiness of malicious BGP hijacks. In *Proc. NDSS* (Feb. 2015).

[26] WAHLISCH, M., SCHMIDT, R., SCHMIDT, T. C., MAENNEL, O., UHLIG, S., AND TYSON, G. RiPKI: The tragic story of RPKI deployment in the Web ecosystem. In *Proc. ACM HotNets* (Nov. 2015).

[27] WHITE, R. Securing BGP through secure origin BGP. *The Internet Protocol Journal 6*, 3 (Sep. 2003), 15–22.

[28] ZHANG, Z., ZHANG, Y., HU, Y. C., MAO, Z. M., AND BUSH, R. iSPY: Detecting IP prefix hijacking on my own. *ACM CCR 38*, 4 (Aug. 2008), 327–338.

[29] ZHAO, X., PEI, D., WANG, L., MASSEY, D., MANKIN, A., WU, S. F., AND ZHANG, L. An analysis of BGP multiple origin AS (MOAS) conflicts. In *Proc. IMW* (Nov. 2001).

[30] ZHENG, C., JI, L., PEI, D., WANG, J., AND FRANCIS, P. A lightweight distributed scheme for detecting IP prefix hijacks in real-time. In *Proc. ACM SIGCOMM* (Aug. 2007).