Studying Impacts of Prefix Interception Attack by Exploring BGP AS-PATH Prepending

Ying Zhang
Ericsson Research, San Jose, USA

Makan Pourzandi
Ericsson Research, Montreal, Canada

***Abstract*—The AS path prepending approach in BGP is commonly used to perform inter-domain traffic engineering, such as inbound traffic load-balancing for multi-homed ASes. It artificially increases the length of the AS level path in BGP announcements by inserting its local AS number multiple times into outgoing announcements. In this work, we study how the AS path prepending mechanism can be exploited to launch a BGP prefix interception attack. Our work is motivated by a recent routing anomaly related to AS Path prepending behavior, *i.e.,* Facebook's traffic being redirected to Korea and China due to a shorter path with fewer prepending ASNs. In order to measure the possible impact of the attack, we develop a simulator to quantify the damage of the attack under a diverse set of attacker/victim combinations. Our main contribution is to quantify how many ASes may be susceptible to the attack, and analyze how effective the attack may be through simulation. Furthermore, we propose an algorithm to detect the interception attack by exploiting inconsistencies via collaborative monitoring from multiple vantage points. Our evaluation shows up to 99% accuracy with 150 vantage points.**

I. Introduction

With a rapidly growing number of critical applications deployed on the Internet today, Internet security has increasingly become an area of concern. The Border Gateway Protocol (BGP) [1] is the de-facto inter domain routing protocol that Autonomous Systems (ASes) use to exchange routes to reach destination address blocks (or *prefixes*). It plays an important role in the well-being of Internet infrastructure. However, due to the lack of security in its original design, the interdomain routing protocol BGP is subject to several types of mis-configuration and attacks [2], [3]. One type of routing protocol attacks is the prefix hijacking attack, which injects and propagates false routes to the Internet, potentially causing traffic to be redirected to the attacker networks [4], [5], [6]. Two general types of prefix hijacks are: injecting a bogus route with a false origin AS (*i.e.,* origin AS attacks) and injecting an incorrect route with a false AS path segment but a legitimate origin AS (*i.e.,* interception attack). The origin AS hijacking attack changes the owner of the prefixes, which usually will cause traffic blackholing [7], [8], [9]. The latter type of attacks is less obvious to the end user as the attacker hijacks the traffic to a prefix and then still forward this traffic on to the prefix owner [5], [10]. Instead of blackholing the destination's traffic, it allows the AS to intercept the traffic without disrupting the destination's reachability.

A special case of the second type of attacks proposed by Ballani *et al.* [5] has demonstrated that an attacker can transparently intercept the hijacked traffic by forwarding it to the prefix owner. It can be achieved by dropping ASes in the AS-PATH attribute in the BGP announcement to make the AS path arbitrarily shorter. For instance, the attacker M modifies the original route to reach prefix owned by V, $ABCV$ to MV, which is much shorter and thus will be adopted by most ASes on the Internet. However, M and V may not have links between them in reality, thus, this type of attack introduces non-existing links. Therefore, such attack can be detected by monitoring anomalous presence of links on the Internet [11].

In this paper, we study a new type of similar attack and show how an attacker can create an interception attack without introducing any anomalous links. This attack exploits a traffic engineering practice: the victim AS V prepends its ASNs multiple times in the AS-PATH attribute. When the attacker M receives the route with k prepended ASNs, it modifies the route intentionally to remove the $k-1$ V's prepended ASNs. As the modified route is much shorter than other routes, most ASes on the Internet would select the route traversing M. The attacker can then eavesdrop on the re-routed traffic, throttle it or even in some cases modify it. Note that the traffic will eventually reaches the destination V, which makes this attack different from the blackholing based prefix hijacking attacks. Different from the previously proposed interception attacks [5], [10], the ASPP-based interception does not remove any real links on the AS path and thus does not introduce any anomalous AS-level links which does not exist in reality.

In this paper, we present a special type of interception prefix hijacking, by exploiting the AS-Path prepending features in BGP protocol, called *"ASPP-based interception attack"*. More specifically, we show how the attacker can achieve the attack without changing or removing any real links between two ASes in the AS path.

Our main contribution is to conduct a comprehensive analysis on the feasibility and effectiveness of ASPP based interception attack. We first characterize the ASPP behavior on the Internet globally from the measurement of public

BGP data over a long duration. We develop a simulator to simulate the impact of such attacks, while considering the AS business relationship induced local preferences in route selection. We found that up to 30% of routes have prepended AS paths, indicating that ASPP is a common practice on the current Internet routing structure. We then develop a methodology to analyze the attack's effectiveness using simulation on the real Internet topology. In particular, we investigate what type of ASes are likely to be hijacked, how the range of polluted ASes is related to the locations of the attacker and victims, and how the number of paddings affects the impact of the attack.

Another contribution of this paper is to propose a detection algorithm for the ASPP-based interception attack. The blackholing based ownership hijacking can be detected using either routing plane anomaly detection [7] or data plane active probing [6], [9], [8]. However, the previous approaches are not directly applied to detect the ASPP-based interception attack. In this work, we propose a detection algorithm based on passively received BGP routing data from public route monitors. The algorithm detects the attack by searching for inconsistencies between announcements from the attacker to different neighbors with some limitations which we discuss in detail. We further validate the algorithm using both simulation and a real-world example.

The paper is organized as follows. In Section II, we describe the AS path prepending approach, the attack and our analysis methodology. A real-world routing anomaly is described in Section III. The impact analysis is introduced in Section IV and the detection is presented in Section V. We present our measurement and analysis results in Section VI. Related work is given in Section VII and Section VIII concludes this paper with future work.

II. ASPP BASED PREFIX INTERCEPTION ATTACKS

In this section, we formally introduce the concept of AS path prepending and the prefix interception attack exploiting the BGP AS path prepending feature.

A. BGP AS path prepending

BGP is the de-facto inter-domain routing protocol for the current Internet. It is used for adjacent routers or ASes to disseminate routing information across different administrative domains. Each BGP announcement contains a set of attributes. One of the important mandatory attributes is the *AS-PATH*, recording the sequence of ASes through which the message has passed. As an announcement passes, each AS adds (prepends) its own AS number to the front of the *AS-PATH* attribute. BGP is based on path vector algorithm, meaning that the route with shorter paths is generally preferred.

BGP is also a policy-based routing protocol. The network operators can configure BGP in certain ways to influence the route selection both locally and globally. AS path prepending (ASPP) is one of these traffic engineering approaches. Instead of prepending its ASN once to the path, an AS adds its own AS number multiple times to artificially increase the length of the AS path. For instance, a BGP announcement with an *AS-PATH* attribute of$\{AS_1, AS_2AS_2AS_2*\ldots AS_k\}$, where * stands for one or more occurrence of AS_2, AS_2 prepends its own ASN 3 times. The longer the AS path is announced to the EBGP neighbor, the less likely the route will be adopted as the best route by other ASes, indicating that the less incoming traffic will be received from that neighbor. When manipulating AS paths, the only valid AS number to prepend is the AS number of the sender. Prepending any other's AS number is considered as misbehaving. The AS which prepends its AS number multiple times into the AS path is called *prepending AS* or *padding AS*.

Through ASPP, an AS could influence the route selection process and thus affect the distribution of traffic flowing into it. However, the AS path length is not the only metric affecting route selection. When multiple routes are available, BGP first selects the route based on local routing policy, which has a higher priority in the decision process than the AS path length. On the other hand, ASPP can be more powerful than other BGP attributes like "Local Preference" when there is a large difference in upstream ISPs' connectivity to the remaining Internet, due to its global impact. Our analysis takes the effect of routing policy into consideration.

The ASPP can be classified into two types, source prepending and intermediary prepending, based on the location of the prepending AS. Source prepending is referred to the case that padding is performed by the origin ASes or the owner of the prefixes, while the intermediary prepending is performed by other non-origin ASes along the path. The AS-path prepending can be configured on the routers using route-map with the set as-path prepend command. It is commonly supported in today's commercial routers.

B. Definition of ASPP-based interception attack

In this section, we detail ASPP-based interception attack exploiting and its impact.

1) Definition of the attack: IP prefix hijacking occurs when a misconfigured or malicious BGP router in a network M either originates or announces a bogus route that traverses M for an IP prefix p owned by another network V. Due to the lack of widely deployed security mechanisms to ensure the correctness of BGP routing updates, the bogus route may be adopted and propagated by some other networks, causing their forwarding tables being polluted. As a result, some of the traffic destined to the victim prefix p is misrouted to the attacker BGP router in M, which can perform any malicious activities. One type of attack is the prefix ownership hijacking or origin AS hijacking, where M acts as the owner of prefix p and modifies the route from $[\ldots V]$ to $[\ldots M]$.

As a consequence, the traffic from polluted ASes are sent to M instead of V, resulting blackholing to V. However, this attack usually triggers the MOAS (Multiple Origin AS) anomalies, which can be easily captured by routing anomaly detection tools [7], [8], [9].

Alternatively, M can advertise the prefix p with AS-PATH $[MV]$, keeping the origin AS the same. This type of attack is referred to as prefix interception or invalid next hop attack [5]. Upon receiving the hijacked traffic to p, M forwards it to V along its existing valid route before hijacking. It allows the attacker to intercept the traffic without disrupting the destination's connectivity to the Internet. However, it may introduce non-existent AS level edges, such as link between M and V. It can be detected as an anomaly of topology changes [11]. Moreover, M should carefully select whom to announce to, to ensure its own valid route to the origin AS V is not affected.

The ASPP-based interception attack is though different as it intercepts traffic without introducing either MOAS or non-existent link anomalies. It does not drop any AS from the original AS path. This type of prefix interception attack exploits the feature of AS-Path prepending mechanism, called *"ASPP-based interception attack"*. In this hijacking, the attacker M receives the original route in the form of $[*VV\dots V]$, where "*" stands for one or more other ASes traversed before M. In the original route, the victim AS prepends its own ASN V multiple times to make it less preferred, for traffic engineering purpose. The attacker modifies the route by removing the duplicated Vs and sends out the malicious route in the form of $[M*V]$. Since the malicious route is shorter than it should be, it is more likely to be chosen as the best route and propagates. Each network either receives the bogus route or may not at all observe the bogus route. In the former case, an AS may choose the bogus route in case the route is more preferred and thus becomes polluted. Once an AS is polluted, all its traffic sent to V will traverse attacker's network M. The attacker can then eavesdrop on the re-routed traffic, throttle it or even in some cases modify it. Note that the traffic will eventually reaches the destination V, which is very different from the blackholing based prefix hijacking attacks.

Note that the prepending is not limited to the origin AS. It can be any ASes who perform AS path prepending before the attacker.

2) Impact of the attack: The hijacking process is as follows. The origin AS V announces prefix p with λ copy of its own ASN, $r_0 = [\underbrace{V\dots V}_{\lambda}]$. The route is sent to V's neighbors and then propagated to other ASes. The hijacking AS M receives this route after traversing n ASes $r_1 = [AS_n \dots AS_1 \underbrace{V\dots V}_{\lambda}]$. AS M removes $\lambda-1$ instances of V at the end of the path and then propagates the route $r_2 = [MAS_n \dots AS_1 V]$. Therefore, r_2 becomes $\lambda - 1$ shorter than r_1. We evaluate the impact of the attack on any AS X by examining if X will choose this invalid route r_2 with one V over the original route with λ number of Vs. If X chooses r_2, M can intercept any traffic originating at X to p, plus any traffic sourced from X's neighbors. AS X's choice depends both on the length of AS path and its routing policies.

Hop	Delay	IP	ASN
1	1 ms	192.168.1.1	
2	41 ms	70.130.143.24	AS7132
3	41 ms	151.164.14.131	AS7132
4	41 ms	151.164.102.106	AS7132
5	131 ms	12.123.30.133	AS7018
6	131 ms	218.30.54.169	AS4134
7	132 ms	202.97.50.37	AS4134
8	137 ms	202.97.49.206	AS4134
9	224 ms	218.30.54.78	AS9318
11	224 ms	198.32.176.71	AS9318
12	245 ms	74.119.77.128	AS32934
13	248 ms	204.15.20.51	AS32934
14	249 ms	69.171.224.39	AS32934

Table I
TRACEROUTE FROM US TO FACEBOOK DURING THE ANOMALY INSTANCE

However, quantifying the impact of such attack is challenging. A common local preference policy is that, for the same destination prefix, an AS prefers to send traffic through a customer than a peer, and it prefers to use a peer than a provider [12]. This is because an AS does not need to pay for the traffic going through its customer link but has to pay for traffic traversing the provider links. In summary, asserting the impact of any ASPP-based interception is not quite straightforward.

III. THE FACEBOOK ROUTING ANOMALY INSTANCE

The Facebook anomaly case is detailed below and used to illustrate the ASPP-based interception attack. The abnormal forwarding behavior for traffic destined to the popular social network site Facebook has been reported in [13]. The AT&T was routing traffic to Facebook through Korea and China (China telecom AS4134).

As mentioned before, ASPP is commonly used for local load balancing and for backup route provisioning. Previous measurement study showed that 32% of the routes in the AT&T network have some form of prepending [14].

Facebook announces the route with 5 duplication of its own ASNs, *e.g., 7018 3356 32934 32934 32934 32934 32934*. The anomaly occurred when a shorter AS path was announced via the Korean ISP AS9318, *i.e., (7018 4134 9318 32934 32934 32934)*, with only 3 padding ASN 32924. This change may be caused by Facebook's misconfiguration or other ASes purposely remove the prepended ASNs 32934. The latter case is an example of an interception attack exploiting the feature of AS path prepending. Although the cause is still not clear to the public, this instance motivates

our work to conduct a systematic study on the ASPP-based interception attack.

More precisely, we gathered the routing tables and updates from route monitors in RouteView [15] and RIPE [16] on Mar. 22nd, 2011. By looking into all the updates associated with any prefixes announced by Facebook (AS32934), we observed that the anomalous route, *(4134 9318 32934 32934 32934)*, appeared at 7:15:02 GMT and is adopted by almost all large ISPs. For example, we observed that AT&T (AS7018) changes its route to Facebook via *(7018 4134 9318 32934 32934 32934)*. The problem is not limited to AT&T, another Tier-1 AS, NTT also chose the same route *(2914 4134 9318 32934 32934 32934)*. However, among all ten prefixes announced by Facebook, only two prefixes, *69.171.224.0/20* and *69.171.255.0/24*, are affected. Using Planetlab based traceroute experiments, we found that most of the Facebook front-end web servers are in these two prefixes. This cross-ocean detour is further verified by traceroute shown in Table I. The traceroute is conducted from a customer of AT&T accessing Facebook. It is shown that the data path is consistent with the BGP routing path, experiencing longer delay than usual.

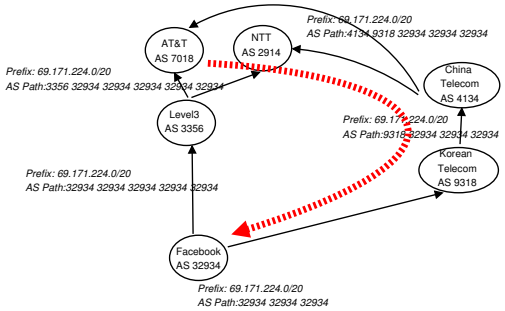

Figure 1. Facebook routing anomaly instance at 7:15am Mar. 22nd

To better understand why this anomalous route is chosen, we retrieved the old route used by AT&T from the routing table snapshot before the change. The BGP level changes are shown in Figure 1. The red dash line shows the flow of the traffic. Interestingly, the old route contains 7 hops, *7018 3356 32934 32934 32934 32934 32934*, which is indeed worse than the anomalous route through China. We further examined the routing tables from previous five days and the next five days after the anomaly. All stable routing tables contain this 7-hop route, which indicates that this is the normal route to reach Facebook from most of the Internet.

To summarize, in the normal case AT&T traverses through Level3 (AS3356) to reach Facebook directly from a 6-hop route. Please note that among the 6 ASNs in the route, 5 of them are the Facebook's AS number 32934. At 7:15 on Mar. 22nd, a shorter route is announced from Korean ISP SK Telecom and then through China Telecom with 5 hops. The 5-hop route only contains 3 Facebook's ASN 32934. From the traceroute we observe that the 5-hop route results in much longer delay, thus considered as an instance of routing anomaly.

There are many possible scenarios explaining the above anomaly. From most monitoring vantage points in US, it is hard to determine which one is the actual cause. The first likely cause is that Facebook purposely announced a shorter route to Korean ISP for traffic engineering by only prepending its ASN twice in the route. This is the most straightforward explanation since AS prepending is commonly used to influence other ASes' routing decision. The second likely cause can be that the Korean ISP AS9318 removed two of 5 prepended Facebook ASNs and sent the new route to its peer China Telecom. China Telecom just prepended its own ASN 4143 and then announced to the rest of the world. The third possibility is that China Telecom modified the route directly by removing two of Facebook's ASNs. In either cases, whether intentionally or by error, the last two cases can be used to illustrate *prefix hijacking* attacks.

IV. Impact Analysis Methodology

In this work, we quantify the impact of ASPP-based prefix interception using simulations on inferred AS topologies.

A. Data source and preprocessing

Our study draws on BGP update messages from the publicly-available RouteViews [15] and RIPE NCC [16] servers. These public servers collect update messages by establishing eBGP sessions with routers in participating ASes. The logs contain the best route from all the peering routers. Our study uses routing table and update data from 2010 to 2011 because the set of monitors are consistent over this period. We start with an initial BGP routing table and apply the stream of update messages to construct a complete view of the latest AS topology.

We combine routing table files in the most recent three months to identify the AS level links. We then infer the AS relationship using the topology graph. We first generate graphs using Gao's algorithm with only Tier-1 peering links as the initial input [17]. We did the same calculation using CAIDA's algorithm. Then we take the set of relationship pairs upon which both graphs agree. We take the common set as the new initial input to re-run Gao's algorithm to generate our topology graph. Although obtaining a precise AS topology is difficult, we attempt to improve its accuracy to the maximum extend.

B. Hijacking simulation

We develop a simulator to emulates BGP route update propagation and the BGP decision process. The routing policies are configured at each AS based on AS relationships. The route selection process follows the "valley free" profit-driven policy. AS path "valley-free" policy permits AS paths in the form of *Customer-Provider* Peer-Peer?*

1. V prepends its ASN λ time in route r
2. Compute all *uphill* paths from V's to all ASes,
path is stored in array PD and distance is stored in D
3. *function* sim_ASPP_hijack(k,M,D,PD)
#$D = \{d_{kj}\}$: distance from ASk to nei j, PD: array of paths for D
if $\exists j, D_{kj} < \infty$ # choose customers' path
if ASk=M
change path $[M * V \dots V]$ to $[M * V]$
store path in PD, return D; update shortest uphill paths;
else # find paths through peer or provider
if $\exists j \in ASk'speers, D_{kj} < \infty$ # choose peers' path
if ASk=M
change path $[M * V \dots V]$ to $[M * V]$
store path in PD, return D;
update shortest path with one peering link;
else # choose provider's path
for each of k's provider p, sim_ASPP_hijack(p,M,D,PD)
update each provider recursively

Figure 2. BGP route update propagation and decision process simulation algorithm

*Provider-Customer**, where "*" represents zero or more occurrence of such type of AS edge and "?" represents at most one occurrence. Thus, in the route selection process, customer's route is preferred over peers route, and peer's route is preferred over the provider's route. This policy is often expressed in the BGP local preference field. With a tie policy-wise, the shorter route is preferred.

The simulation algorithm is shown in Figure 2. The victim first announces its route with λ times of prepending (step 1). According to the latest AS-level path simulation algorithm [18], we first compute the shortest uphill paths from V to all other ASes (step 2). The uphill path only contains the customer-provider links (up links). If there is no customer route, we then search peering routes and finally provider routes. In each step, if the current AS is the attacker M, it removes $\lambda - 1$ Vs, keeping only one copy of V's ASN in the route. Because this route is shorter, it may affect the previous shortest uphill path selection, we need to update the shortest uphill paths accordingly.

For each experiment, we fix a victim and an attacker AS to simulate the interception. We quantify the impact of the attack by the fraction of ASes adopting the malicious route, meaning that their traffic to victim V will traverse attacker M. In later Section VI, we will present results of different levels of impacts under various combinations of V and M.

V. Detection of ASPP based prefix interception attack

Our basic approach is to examine BGP routing data collected by the route monitors, for example RouteViews, RIPE, and any other BGP collectors, and provide real time notifications of any potential ASPP based prefix interception hijacking to the prefix owner in a reliable way. In particular, we should raise alarms for the prefix owner who is performing AS path prepending when a malicious AS modifies the prepended ASNs in the AS path to create the interception. In practice, an prefix owner can monitor the data from public monitors continuously using tools like PHAS [7].

A. Algorithm design

The main challenge in detection is that the origin AS can apply flexible prepending policies for traffic engineering purposes. For instance, the origin AS may send a legitimate shorter route (*i.e.,* with fewer prepended ASNs) to a particular neighbor so that more traffic will traverse this neighbor. In another extreme case, for the purpose of provisioning backup routes, the prefix owner can choose the degree of padding to be large enough so that the backup route will not be adopted as best route unless if there is a failure in other primary routes. In order to lower false positives, the detection algorithm must differentiate the malicious case from other legitimate reasons for changing prepending behaviors.

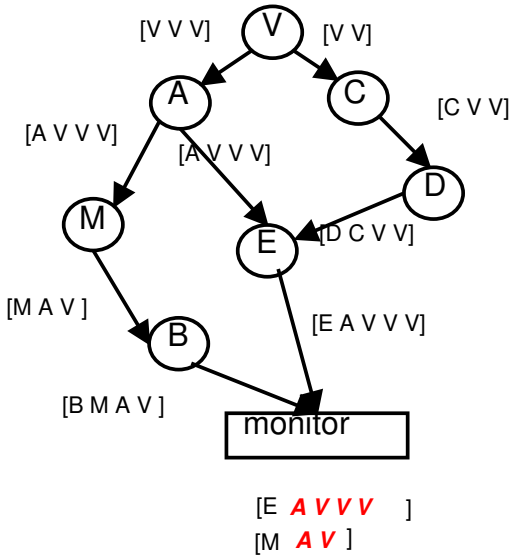

Figure 3. Example of ASPP based prefix injection attack detection

The detection is based on the observation of BGP route propagation: following the same AS path, at any given time, an AS cannot receive two routes with two different padded ASNs. A brief explanation is provided below. Given two routes $r_0 = [AS_k AS_{k-1} \dots AS_1 \underbrace{V \dots V}_{\lambda}]$ and $r_1 = [AS'_k AS'_{k-1} \dots AS'_1 \underbrace{V \dots V}_{\lambda'}]$, let's consider the scenario of $[AS_k AS_{k-1} \dots AS_1] = [AS'_k AS'_{k-1} \dots AS'_1]$. Assuming each AS has consistent routing policy to the same neighbor [12], from r_0 and r_1 we can infer that the victim sends $[\underbrace{V \dots V}_{\lambda}]$ and $[\underbrace{V \dots V}_{\lambda'}]$ to its neighbor AS_1. Since an AS should apply the same policy to the same neighbor, we have $\lambda = \lambda'$.

We use an example in Figure 3 to illustrate the detection process. In this example, victim V announces two routes

function $DETECT\_ASPP\_INTERCEPT(r_t^d, r_{t-1}^d)$
#r_t: current route from AS d to reach prefix p;
#r_{t-1}^d: last route before changes
$r_t^d = [AS_I AS_{I-1} \dots AS_1 \underbrace{V \dots V}_{\lambda_t}]$# current route
$r_{t-1}^d = [AS_J AS_{J-1} \dots AS_1 \underbrace{V \dots V}_{\lambda_{t-1}}]$# previous route
if ($\lambda_t < \lambda_{t-1}$) # padded number decreases
  Search routes from all n ASes to p at t: $R = [r_t^1, r_t^2, \dots r_t^n]$
  for each ($r \in R$)
    $r = [AS'_L, AS'_{L-1} \dots AS'_1 \underbrace{V \dots V}_{\lambda_l}]$
    if ($[AS'_{L-1} \dots AS'_1] = [AS_{I-1} \dots AS_1]$)
      if ($\lambda_t < \lambda_l$)# route from AS_I has fewer padding
        Raise Alarm: *detect attack!*
        AS_I removes $\lambda_l - \lambda_t$ padded ASNs
    else # no direct symptom, search for possible hints
      if ($\lambda_t < \lambda_l$)
        if ((length(AS'_L to V)+λ_l)>(length(AS_I to V)+λ_t))
          if (AS_{I-1} is customer of AS'_L)
            Raise Alarm: *possible attack!*
          else if ((AS_{I-1} is peer of AS'_L) AND (no peer links in r_t^d))
            Raise Alarm: *possible attack!*
          else if ((AS_{I-1} is provider of AS'_L) AND
          (AS'_{L-1} is a provider of AS'_L))
            Raise Alarm: *possible attack!*

Figure 4. ASPP based prefix injection detection algorithm

with 2 pads ($[VV]$) and 3 pads ($[VVV]$) to its neighbors A and C respectively. The purpose could be victim V wants more traffic traversing C as it has lower cost. Each other AS adds its own ASN once to the beginning of the AS path. Besides prepending, Attacker M removes the 2 Vs from the AS path and sends a bogus route $[MAV]$ to its neighbor B. In the monitor, we observe two conflicting path segments $[EAVVV]$ and $[MAV]$. It is impossible for V to send two different routes to A, as we assume the same policy is applied to the same neighbor. Moreover, A does not have incentives to send two routes with different paddings of V to M and E. If A wants to perform traffic engineering, it would prepend its own ASN A multiple times. Therefore, the algorithm detects that M most likely modifies the route.

Figure 4 shows the details of the algorithm. For each routing change to a shorter AS-path due to fewer padded ASNs from AS d, we search for routes to the same prefix from all other ASes by combining views from all monitors. Please note that the total ASes n are larger than the number of monitors, as destination based routing. Among the set of all routes observed R, we search for any path containing the same path segment as the current route from the second AS AS_{i-1} to the second last AS AS_1. Once found the route with common path segment, if the padded ASN instances are different, more precisely if the current route contains fewer pads, then we detect an inconsistency between these two routes. More specifically, the same sequence of ASes $AS_{i-1} \dots AS_1$ receives two routes with different padding V set, $\underbrace{V \dots V}_{\lambda_t}$ and $\underbrace{V \dots V}_{\lambda_l}$, which results in inconsistency. Therefore, it suggests that AS_i modifies the route r_t^d by removing padded Vs to make it shorter.

However, there is no guarantee that we can always find a common segment from AS_{i-1} to AS_1 with different paddings. If the symptom is found, we raise an alarm that the attack is detected with high confidence. If not, we continue searching for hints for the attack. If observing another AS'_L, a neighbor of AS_{i-1}, selects a longer route with λ_l paddings, it is only possible if AS_{i-1} does not propagate the shorter route to AS'_L. We examine if it is possible given the route selection policies and their AS relationship. If there is local policies preventing AS'_L receiving the shorter route from AS_{i-1}, we raise an alarm with lower confidence, given that the inferred AS level relationship might be inaccurate.

B. Limitations

The basic idea of the detection mechanism is essentially to search for inconsistent route advertisements as an indicator of malicious modification of an advertised route, which has been used in many detection proposals for different routing attacks [19], [7], [20]. Our contribution in this work is mainly building our method on top of this general concept and making it specific to detect the ASPP based interception attack.

Overall, detection on the ASPP-based interception is difficult as only limited information is available to the public. Similar conclusions have been also drawn for invalid next hop based interception attack [5]. Our detection takes a two-step approach, *i.e.,* drawing conclusions with different confidence depending on the available data.

Ideally, if the prefix owner has monitoring vantage points in all the ASes, then the detection accuracy is 100%. However, in reality, the detection is affected by the distribution of vantage points. There are a few corner cases that the detection may fail. First, if a direct neighbor of the victim is the attacker, then the victim cannot detect the attack unless it has a vantage point on the attacker, or on any of the attacker's direct neighbors. With routes monitored from any of these vantage points, the victim can detect an inconsistency between the actual route announced and the modified route. Second, the location of the vantage points limits the range of attacks that can be detected. In this case, each victim can select a set of important ASes as their monitors to prevent being hijacked. In our future work, we will study the selection of vantage point to perform self-defense for different victims.

In summary, we admit that it is not guaranteed to detect the attack. However, it is a practical solution in reality given not all ISPs are willing to share their routing data.

VI. Experimental results

Below we present our results by analyzing the BGP data as well as simulations. We first start with a characterization of ASPP behaviors from BGP routing data collected from

RouteView and RIPE. It shows how common it is used on the Internet. These results show that there exists many opportunities for ASPP based interception attack to occur on the Internet today. Then we examine the effectiveness of prefix hijacking attack from multiple aspects through our simulator.

A. Usage of AS Path Prepending on the Internet

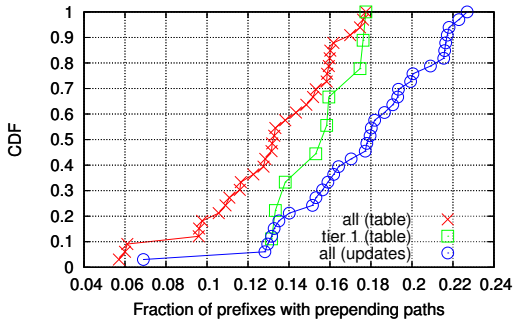

Figure 5. Fraction of routes with prepending ASes

It is known that ASPP is commonly used for a non trivial portion of prefixes by most ASes [21]. We confirm this observation using more recent data in Mar. 2011. Figure 5 shows the fraction of prefixes whose best route contains ASPP, compared to the total number of prefixes. We calculate one fraction number for each monitor and present the CDF of many monitors. On average 13% routes in the default routing table have ASPP. Among different monitors, we observe a significant difference. We suspect that edge ISPs are less likely to see prepending paths because the prepending paths are longer and thus are less likely to be selected as the best route. The top tier ISPs, on the other hand, may observe a larger fraction, since they observe a diverse set of routes given their larger base. This conjecture is confirmed when we also plot the fractions for only Tier-1 ISPs in Figure 5. For similar reasons, in the update files, we also observe more routes with prepending ASes. In the unstable states, these routes are more likely to be visible in the route monitoring system. Overall, these results show that ASPP is a commonly used practice on the Internet. Therefore, ASPP based interception attack is likely to happen on the Internet today.

The operators configure the ASNs to be duplicated different number of times for different preference. For a less preferred route, it may repeat it many times to ensure it won't be chosen as the best route. We study how many repetition is common in most prepending AS paths. Figure 6 shows the number of duplicated ASNs in all the routes. Most of them are very small: 34% repeat twice and 22% repeat three times observed from routing table. 1% of them repeat larger than 10 times. The routes from update files have larger duplications. Please note that many routes that have a higher number of prependings were most likely not selected by the transit ASes and therefore filtered out. Unless we analyze the data from more vantage points, we cannot accurately estimate the number of such prepended routes. The large number of prepended instances provides opportunities for the attacker to remove some of them, making the AS path shorter.

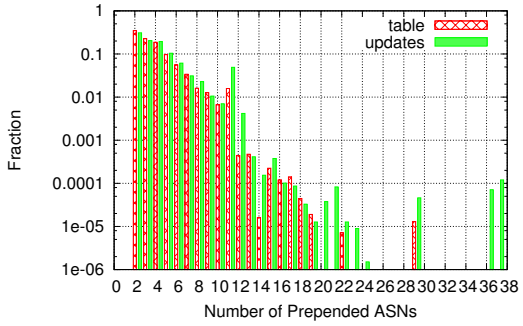

Figure 6. Number of duplicate ASNs

B. Impact analysis of ASPP-based prefix interception attack

Below, we study the impact of an ASPP-based interception attack from the following perspectives. First, the location of the attacker and victim has large impact on the attack effectiveness. We further illustrate the impact using a few special examples of attacker/victim combinations.

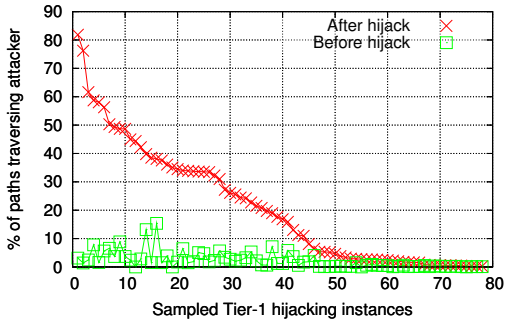

Figure 7. Polluted ASes in attacks between Tier-1 ASes (prepended ASN=3)

1) Attacker and victim locations: We first study that if the hijacking AS is a Tier-1 AS, what fraction of other Tier-1 ASes can be intercepted. A tier-1 AS is an AS with no providers and is peering with all other tier-1 ASes. We treat it as a special case since Tier-1 AS is likely to be traversed by many paths given its core position on the Internet. In such cases, each AS will receive the invalid route and the original route either through a provider link or through a peering link. An AS is not polluted only if it is a direct or indirect customer of the victim or it is a peer of victim's customers. With a Tier-1 AS being the attacker, we further

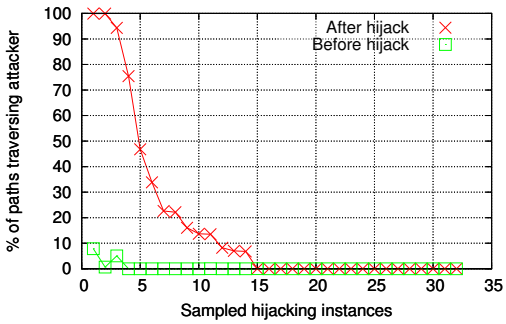

Figure 8. Polluted ASes in attacks between randomly sampled ASes

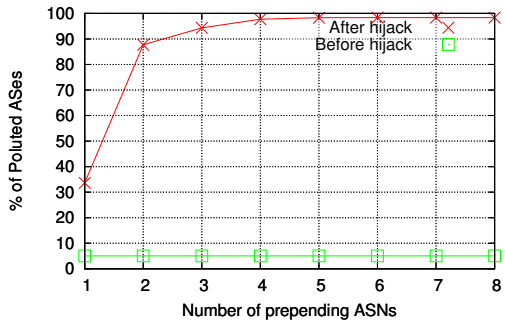

Figure 9. Pollution range with diff. prepended ASNs (AS1239 hijacks AS7018)

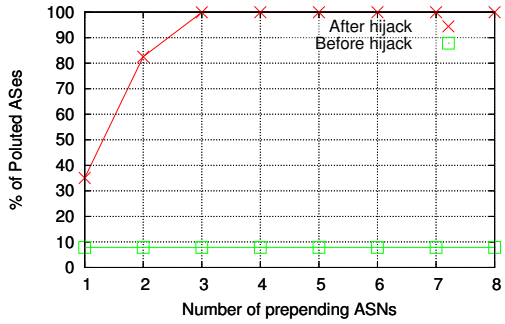

Figure 10. Pollution range and prepended ASNs (AS7018 hijacks AS32934)

study two sub-categories, the victim also being Tier-1 AS or the victim being a lower Tier AS. Intuitively, victims closer to the core of the Internet would have more resilience to attacks, as it has on average shorter hops to reach the rest of Internet.

Figure 7 shows the ASPP interception where both the attacker and the victim are Tier-1 ISPs, which are in the same topological hierarchy and with similar connectivity characteristics. Moreover, being in the core of the Internet, the attack can intercept a large portion of Internet traffic. We simulate 80 instances of such hijacking cases with 3 prepended instances. We choose 3 ASNs to pad because it is half of the average AS path length. Figure 7 presents the results for all hijacks ranking based on the pollution range. Overall, the pollution range is around 40%, not as high as we expected. Interestingly, the last 30 hijacking cases have only less than 5% impact. With more investigation, we found that in these cases, the victim's customers are richly peered with other ASes, effectively spreading out the legitimate routes.

Besides Tier-1 ASes, we also conduct experiments on randomly selected ASes. In these experiments, both the attackers and the victims are randomly selected, most of which are Tier-4 and Tier-5 ASes. Figure 8 shows the 27 experiments ranked by the fraction of polluted ASes. Interestingly, the hijacks are less effective in most cases compared to the Tier-1 ASes. One reason is that only a small number of paths to the victim will traverse the attacker in the normal case, so that the attacker has only limited opportunities to remove the prepended ASNs. Another factor is that the attacker is at the edge of the Internet, with relatively long path to reach other ASes, even after removing the duplicated ASNs.

2) Special attack scenarios: The number of ASNs prepended by the origin AS directly relates to the effectiveness of the interception attack. The more hops being prepended, the longer the AS path will be, resulting in larger chance for the modified path from the attacker to be preferred. We conducted a set of experiments with different numbers of prepending ASN hop counts. For each hijacking instance, we compute the fraction of paths traversing the attacker, meaning that the hijack succeeds. For the ease of comparison, we also plot the fraction of paths traversing the attacker ASN in the normal scenario without hijacking in the same figure. Figure 9 shows the instance of Sprint (AS1239) hijacking AT&T (AS7018). Since both ASes are large ISPs with diverse connectivity to the rest of the Internet. With only one prepended ASN, only 30% of the ASes on the Internet switch their current paths to AT&T's prefixes, traversing Sprint. However, when the number of prepended ASNs increases to 2, 80% of the ASes choose the hijacked route. As the prepended ASN increases to 3 and 4, more than 95% of the paths switch. When the number of prepended instances reaches 5 and above, the pollution range remains the same. These ASes could be AT&T's single homed customer. It could also be some Tier-2 ASes who are AT&T's direct peers, which would prefer a peering route instead of going through providers.

We next examine the two ASes with completely different topological characteristics, a high tier AS hijacks a lower tier AS, and vise versa. The first example is a Tier-1 attacker AT&T (AS7018) attacks a Tier-3 victim Facebook (AS32934). In this case, the lower tier victim's route will be preferred and adopted by its providers, the providers' providers, and their direct peers. For the remaining ASes, they are more likely to choose the invalid route from the

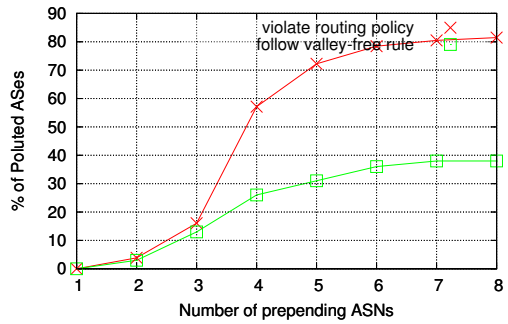

Figure 11. Pollution range and prepended ASNs (AS32934 hijacks AS2914)

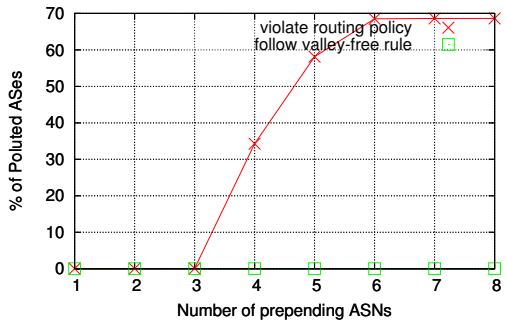

Figure 12. Pollution range and prepended ASNs (AS30209 hijacks AS12734)

higher tier attacker, since both the original route and the invalid route will reach them through provider or peering links. Thus, the impact of the attack depends on how rich the connectivity of victim's providers and peers are. The prepended hops would have larger impact on other ASes receiving both routes from provider links. Figure 10 shows that when the prepending instances is larger than 2, more than 99% of the Internet are polluted. 82% of ASes have switched to the malicious route with 2 hop prepended.

On the other extreme, we let the small AS Facebook attacks a Tier-1 AS (AS2914) in Figure 11. In this case, most other ASes originally use providers' route to reach the victim, except for the victim's peers, including the attacker. Since the attacker learns the route from a "down-hill" path, it can not send it to its providers, according to "valley-free" rules. Therefore, the attacker can only pollute its customers, peers, and peers' customers. Since the attacker is already in the low level of AS hierarchy, it should only have limited impact. However, if the attacker does not obey the "valley-free" rules but sends the route to its other providers, the impact can be equally large as other scenarios, especially with longer padded ASNs. In Figure 11, the polluted range is surprisingly much larger than expected, *i.e.,* around 38% when the number of paddings is sufficiently large. With further investigation, we found that the AS2914 is a sibling of popular CDN Limelight, which happens to be a customer of Facebook. After hijacks, Facebook sends the invalid route to its provider Akamai, which peers with other 235 ASes. The entire process obeys the "valley free" routing policy. Since many of Akamai's peers are large ISPs, the bogus route is propagated widely. This example illustrates that Internet is better connected than expected, *i.e.,* a small but well-connected enterprise ISP can even intercept a Tier-1 ISP's traffic, under certain AS path prepending conditions.

Besides Tier-1 ISP as attacker or victim, we examine smaller ASes for both attacker and victim in Figure 12. It is similar to the previous scenario in Figure 11. If the attack does not send the route learned from one provider to another, the set of polluted ASes is very small. However, if it does not follow the "valley-free" rule, the impact can be significant when the victim adds more prependings.

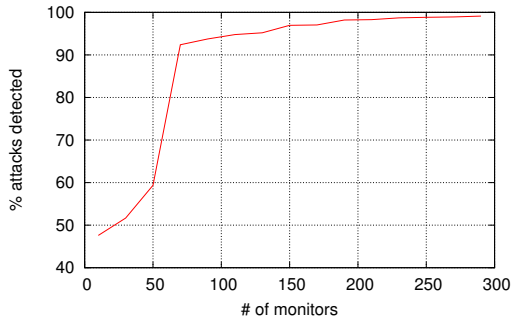

Figure 13. Detection accuracy with increasing monitors

C. Evaluation of detection algorithm

Finally, we evaluate our detection methodology presented in Section V through simulation. For a given set of monitors, we randomly choose 200 attacker and victim pairs to simulate the hijacking process using methods presented in Section IV-B. We examine the number of attacks detected using our detection algorithm. Obviously the capability of detection is affected by the set of monitors. The more monitors there are, and the more diverse they are located, the higher is the accuracy. The selection of monitor sets is a complex problem, which will be studied in our future work. In this validation, we simply rank all ASes based on their degrees and select the top d monitors. Figure 13 shows the percentage of detected attacks as d increases. We observe that with 70 monitors we can detect 92% of the attacks. When the set of monitors is increased beyond 150, the accuracy is above 99%.

Another important metric to evaluate a detection algorithm is how fast it can detect the attack after it is launched. Ideally if we have monitors deployed surrounding the attacker, then the inconsistency can be observed right after the attacker propagates the route. However, due to the limited locations of the monitors, by the time the attack is detected, a set

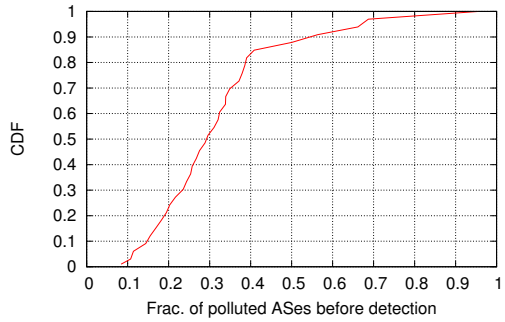

Figure 14. Fraction of ASes polluted before detection

of ASes are already polluted. We quantify this using the fraction of ASes polluted, *i.e.,* adopting the modified route as the best route, before the attack is detected by the monitors. The speed of detection tightly correlates with the monitors' locations as well. We use the top 150 ASes ranked by degrees as the monitors and compute the fraction of polluted ASes for each experiment. We repeat the experiment 200 times with random victim/attacker pairs. Figure 14 shows the CDF of the fraction of polluted ASes. It shows that we can detect the attack in fairly early stage, 80% of experiments with less than 37% polluted ASes.

VII. Related work

There exist several public route monitoring systems, such as Route Views [15] and RIPE [16], to help understand and monitor the Internet routing system. These monitoring systems operate by gathering real-time BGP updates and periodic BGP table snapshots from routers to discover dynamic changes of the global routing system. Various research studies have been conducted relying on these data, including network topology discovery [22], AS relationship inference [23], [24], [25], [26], [12], AS-level path prediction [18], [27], BGP root cause analysis [28], and several routing anomaly detection schemes.

Our work is closely related to previous work on IP prefix hijacking detection and prevention. Previously, researchers have proposed the detection systems in control plane [7], [29], [3] that rely on detecting suspicious MOAS activity. After detection, they either simply alert victim prefix owner [7], or delay the propagation of suspicious routes [29]. Hijacking can also be detected using data plane active measurement [5], [9]. These schemes rely on observing inconsistent data plane network properties or end-host based properties [6].

Upon detecting prefix hijacking events, the natural next step of action is to mitigate their impact. Numerous mitigation schemes have been proposed, including manual response to install filters, ACR [30], MIRO [31], route purge-promotion [32], and overlay routing, *e.g.,* RON [33].

Our detection algorithm relies on passively collected BGP data. It is closely related to the large body of work on passive counter-measures against prefix hijacking. These include cryptography-based approaches such as S-BGP [34], So-BGP [35], SPV [36], listen-whisper [37], and HiBGP [4], and non-cryptography-based approaches such as PG-BGP [29], intentional deaggregation, bogon filter, and customer route filtering [32]. Similar to other proposals relying on BGP data, our detection is inherently limited by the monitor's location [3], [38].

VIII. Discussion and Conclusions

AS Path prepending is a common approach for inter-domain traffic engineering. It relies on manipulating the AS path length by purposely inserting its own ASN multiple times. However, in this work we study a new type of BGP prefix interception attack by exploiting the ASPP mechanism. We present a comprehensive study on the feasibility and effectiveness of such attacks and examine the damage caused by the attack based on three aspects: the location of the victim AS, the location of the attacker AS, and the number of prepended ASNs in the AS path. We developed a simulator to analyze the attack's impact based the real Internet topology.

Furthermore, we propose an algorithm to detect the interception attack by exploiting inconsistencies from multiple vantage points. Our detection algorithm can detect most attacks given a good selection of monitors. However, it is not 100% guaranteed to detect all the attacks. One limitation of our method is that it heavily relies on the location of the monitors. In our future work, we plan to investigate the best vantage point selection to guarantee the detection of the interception attacks. Developing attack prevention schemes is also in our future agenda.

References

[1] Y. Rekhter and T. Li, "A Border Gateway Protocol." RFC 1771, March 1995.

[2] V. J. Bono, "7007 Explanation and Apology." NANOG 97-04.

[3] M. Lad, R. Oliveira, B. Zhang, and L. Zhang, "Understanding resiliency of internet topology against prefix hijack attacks," in *Proc. of DSN*, 2007.

[4] J. Qiu and L. Gao, "Hi-BGP: A Lightweight Hijack-proof Inter-domain Routing Protocol," tech. rep., University of Massachusetts Amherst, 2006.

[5] H. Ballani, P. Francis, and X. Zhang, "A Study of Prefix Hijacking and Interception in the Internet," in *Proc. ACM SIGCOMM*, August 2007.

[6] X. Hu and Z. M. Mao, "Accurate Real-time Identification of IP Prefix Hijacking," in *Proc. of IEEE Security and Privacy*, 2007.

[7] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang, "PHAS: A Prefix Hijack Alert System," in *Proc. of USENIX Security Symposium*, 2006.

[8] Z. Zhang, Y. Zhang, Y. C. Hu, and Z. M. Mao, "iSPY: Detecting IP Prefix Hijacking on My Own," in *Proc. ACM SIGCOMM*, 2008.

[9] C. Zheng, L. Ji, D. Pei, J. Wang, and P. Francis, "A Light-Weight Distributed Scheme for Detecting IP Prefix Hijacks in Realtime," in *Proc. ACM SIGCOMM*, August 2007.

[10] S. Goldberg, S. Halevi, A. D. Jaggard, V. Ramachandran, and R. N. Wright, "Rationality and traffic attraction: incentives for honest path announcements in bgp.," in *Proc. ACM SIGCOMM*, pp. 267–278, 2008.

[11] Y. Zhang, Z. M. Mao, and J. Wang, "A firewall for routers: Protecting against routing misbehavior," in *Proceedings of the 37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, DSN '07, (Washington, DC, USA), pp. 20–29, IEEE Computer Society, 2007.

[12] L. Gao, "On Inferring Autonomous System Relationships in the Internet," in *Proc. IEEE Global Internet Symposium*, 2000.

[13] "Hey ATT customers: Your Facebook data went to China and S. Korea this morning." www.blyon.com/hey-att-customers-your-facebook-data-went- to-china-and-korea-this-morning.

[14] N. Feamster, J. Borkenhagen, and J. Rexford, "Guidelines for interdomain traffic engineering," *SIGCOMM Comput. Commun. Rev.*, vol. 33, pp. 19–30, October 2003.

[15] "University of Oregon Route Views Archive Project." www.routeviews.org.

[16] "Ripe NCC." http://www.ripe.net/ripencc/pub-services/np/ris/.

[17] L. Gao, T. G. Griffin, and J. Rexford, "Inherently safe backup routing with BGP," in *Proc. IEEE INFOCOM*, 2001.

[18] Z. M. Mao, L. Qiu, J. Wang, and Y. Zhang, "On AS-Level Path Inference," in *Proc. ACM SIGMETRICS*, 2005.

[19] L. Subramanian, V. Roth, I. Stoica, S. Shenker, and R. H. Katz, "Listen and Whisper: Security Mechanisms for BGP," in *Proc. first Symposium on Networked Systems Design and Implementation (NSDI)*, 2004.

[20] "Thousand Eyes." http://www.thousandeyes.com/.

[21] J. H. Wang, D. M. Chiu, J. C. S. Lui, and R. K. C. Chang, "Inter-as inbound traffic engineering via aspp," *IEEE Transaction on Network and Service Management*, vol. 4, June 2007.

[22] Y. He, G. Siganos, M. Faloutsos, and S. V. Krishnamurthy, "A systematic framework for unearthing the missing links: Measurements and Impact," in *Proc. of NSDI*, 2007.

[23] X. Dimitropoulos and G. Riley, "Modeling Autonomous System Relationships," in *Proc. of PADS*, 2006.

[24] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, kc claffy, and G. Riley, "AS Relationships: Inference and Validation," *ACM Computer Communication Review*, vol. 37, no. 1, 2007.

[25] G. Battista, M. Patrignani, and M. Pizzonia, "Computing the Types of the Relationships Between Autonomous Systems," in *Proc. IEEE INFOCOM*, March 2003.

[26] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz, "Characterizing the Internet hierarchy from multiple vantage points," in *Proc. IEEE INFOCOM*, 2002.

[27] W. Muhlbauer, A. Feldmann, O. Maennel, M. Roughan, and S. Uhlig, "Building an AS-Topology Model," in *Proc. of ACM SIGCOMM*, 2006.

[28] A. Feldmann, O. Maennel, Z. M. Mao, A. Berger, and B. Maggs, "Locating Internet Routing Instabilities," in *Proc. ACM SIGCOMM*, 2004.

[29] J. Karlin, J. Karlin, S. Forrest, and J. Rexford, "Pretty Good BGP: Improving BGP by Cautiously Adopting Routes," in *Proc. of ICNP*, 2006.

[30] D. Wendlandt, I. Avramopoulos, D. Andersen, and J. Rexford, "Don't Secure Routing Protocols, Secure Data Delivery," in *Proc. of ACM Workshop on Hot Topics in Networks (HotNets)*, 2006.

[31] W. Xu and J. Rexford, "MIRO: multi-path interdomain routing," in *Proc. ACM SIGCOMM*, 2006.

[32] Z. Zhang, Y. Zhang, Y. C. Hu, and Z. M. Mao, " Practical Defenses Against BGP Prefix Hijacking ," 2007.

[33] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris, "Resilient overlay networks," in *Proc. ACM SOSP*, 2001.

[34] S. Kent and C. Lynn and K. Seo, "Secure Border Gateway Protocol (Secure-BGP)," *IEEE J. Selected Areas in Communications*, 2000.

[35] J. Ng, "Extensions to BGP to Support Secure Origin BGP (soBGP)." IETF Draft: draft-ng-sobgp-bgp-extensions-01.txt, November 2002.

[36] Y.-C. Hu, A. Perrig, and M. Sirbu, "SPV: A Secure Path Vector Scheme for Securing BGP," in *Proc. ACM SIGCOMM*, 2004.

[37] L. Subramanian, V. Roth, I. Stoica, S. Shenker, and R. H. Katz, "Listen and Whisper: Security Mechanisms for BGP," in *Symposium on NSDI*, 2004.

[38] Y. Zhang, Z. Zhang, Z. M. Mao, Y. C. Hu, and B. Maggs, " On the Impact of Route Monitor Selection ," in *Proc. ACM SIGCOMM Internet Measurement Conference*, 2007.