

Practical Defenses Against BGP Prefix Hijacking

Zheng Zhang Ying Zhang Y. Charlie Hu Z. Morley Mao
Purdue University University of Michigan Purdue University University of Michigan

ABSTRACT

Prefix hijacking, a misbehavior in which a misconfigured or malicious BGP router originates an IP prefix that the router does not own, is becoming an increasingly serious security problem on the Internet. In this paper, we conduct a first comprehensive study on incrementally deployable mitigation solutions against prefix hijacking. We first propose a novel reactive detection-assisted solution based on the idea of bogus route purging and valid route promotion. Our simulations based on realistic settings show that purging bogus routes at 20 highest-degree ASes reduces the polluted portion of the Internet by a random prefix hijack from 50% down to 24%, and adding promotion further reduces the remaining pollution by 33% ~ 57%. We prove that our proposed route purging and promotion scheme preserve the convergence properties of BGP regardless of the number of promoters. We are the first to demonstrate that detection systems based on a limited number of BGP feeds are subject to detection evasion by hijackers. Motivated the need for proactive defenses to complement reactive mitigation response, we evaluate customer route filtering, a best common practice among large ISPs today, and show its limited effectiveness. We also show the added benefits of combining route purging-promotion with customer route filtering.

1. INTRODUCTION

Internet routing is a critical infrastructure service for distributing reachability information globally. Partly due to the assumption made by the early Internet designers that there exists little or no malicious and misconfiguration behavior on the Internet, today's Internet routing system is still largely unprotected. Unfortunately, we have witnessed several serious incidents [9, 23] of disrupted network connectivity for

many prefixes including those hosting important services such as DNS. Despite many proposals such as So-BGP [25], SBGP [19], and SPV [17], there are still no widely deployed effective prevention and mitigation solutions against routing attacks such as IP prefix hijacking. Two main problems exist with existing solutions, hindering widespread adoption. Firstly, many of these solutions require significant modifications to the BGP routing protocol, making adoption challenging. Secondly, the benefit of partial adoption appears limited, leading to reluctant initial adoption [12].

The critical importance of protecting the Internet from IP prefix hijacking attacks, which can severely disrupt network reachability, motivates the need for devising incrementally deployable network-based defense solutions. Existing work has so far focused mainly on 1) *detection* alone, relying on manual response from network operators, without considering automated responses, and 2) *proactive prevention*.

In this paper, we build on previous work on automatic prefix hijacking detection to propose automatic reactive *mitigation* mechanism in response to detected attacks. Our solution is based on the idea of bogus route purging and valid route promotion. Participating ASes, typically in the core of Internet, delete the bogus routes. Some ASes promote valid routes by shortening their AS paths using the AS_SET construct, while preserving the forwarding path integrity. Based on realistic simulations, we show that with only 20 participating ASes, the percentage of polluted ASes is reduced from 50% to only 15%. Compared to previous work, our scheme can even effectively combat colluding attackers from different network locations. Moreover, we prove that the addition of route promotion does not change the convergence guarantees of the current Internet. We finally study the benefit of incremental deployment in terms of the best placement of mitigation solutions.

In addition to reactive mitigation, we analyze how detection systems relying on multiple BGP feeds are subject to evasion, and demonstrate this limitation using realistic settings. Motivated by the need for proactive prevention to eliminate IP hijacking in many cases to complement reactive mitigation response, we study customer route filtering, a best common practice, and show its limitation. We also show that this proactive scheme can be combined with our

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CoNEXT'07, December 10-13, 2007, New York, NY, U.S.A.
Copyright 2007 ACM 978-1-59593-770-4/07/0012 ...\$5.00.

reactive scheme to provide higher benefit than each of them alone.

The rest of the paper is organized as follows. Section 2 provides background on prefix hijacking and a taxonomy of hijacking defense solutions. Section 3 presents the methodology of our study. Section 4 presents a novel reactive mitigation scheme. Section 5 shows the limitation of the reactive approach due to detection evasion and analyzes a proactive scheme. Finally we conclude with related work and several remarks.

2. BACKGROUND

In this section, we briefly review IP prefix hijacking targeted at the interdomain routing protocol – BGP. IP prefix hijacking occurs when a misconfigured or malicious BGP router in a network N either originates or announces a bogus route that traverses N 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 N , which can perform any malicious activities pretending to be the victim prefix p or may even choose to selectively forward the traffic back to the victim [8]. During a hijacking, the bogus route is of the form $[\dots N]$, whereas the original correct route is of the form $[\dots V]$. Each network M either receives the bogus route or may not at all observe the bogus route. In the former case, M may choose the bogus route in case the route is more preferred and thus becomes *polluted*. In the latter case, M 's neighboring networks must not be polluted, thus preventing M from observing the bogus route.

IP prefix hijacking can be performed in several ways. We describe the two main types to facilitate our subsequent discussion of defense solutions. A more detailed classification can be found in a recent study [16].

1. *Regular prefix hijack* occurs when the attack router originates a route to an existing IP prefix of the victim network. As a result, the Internet is partially polluted, depending on how preferable the bogus route is compared to the valid route from the perspective of various networks.
2. *Subprefix hijack* results from stealing a subnet of an existing prefix in the routing tables by announcing a route for the subnet originating from the attacker network. Due to longest-prefix-matching based forwarding, most networks are polluted.

To increase detection difficulties, stealthy attackers may disguise both attack types with falsified AS paths without modifying the origin AS, while making traffic traverse through the attacker network. Thus, the bogus route will be of the form $[\dots N \dots V]$.

2.1 Taxonomy of Prefix Hijacking Defense

Table 1 presents a taxonomy of the various solutions on defending against BGP prefix hijacking attacks, including detection schemes, and the main existing techniques, as well as two techniques studied in this paper for mitigation and prevention.

There are two main approaches to defending against various security attacks on routing protocols: proactive prevention and reactive mitigation. Ideally, prevention is preferred as it aims to eliminate attacks. However, due to a lack of global adoption of necessary changes required for prevention and the possibility of network misconfiguration, proactive prevention alone is never sufficient. After all, Internet consists of heterogeneous networks, it is quite challenging, if not impossible, to enforce uniformly correct configurations and adoption of any newly proposed changes non-essential to network operations.

It is important to note that reactive mitigation must depend on accurate and timely detection systems to be effective. Besides potential inaccuracies, we demonstrate in Section 5.1 that detection systems relying on multiple BGP feeds from different vantage points are inherently susceptible to detection evasion. Given such limitations, similar to proactive prevention, reactive mitigation is also imperfect. Therefore in this paper we also analyze the effectiveness of a known proactive scheme and the added benefits of combining proactive and reactive approaches. Moreover, we study how deployment locations affect overall effectiveness.

Table 1 further classifies the reactive mitigation into network-based and end-host based schemes. There are clear trade-offs to each category. Network-based detection and response require cooperation from network elements inside the core of the Internet and may suffer from increased route convergence delays. In contrast, an end-host based approach can be more readily deployed by end-users or at the edge of the network, but has more limited scope of effectiveness. It usually relies on application-layer techniques such as overlay routing to bypass polluted networks.

In this work, we focus on incrementally deployable, network-based reactive mitigation and proactive prevention solutions mainly due to their better efficiency and potential for larger scope of impact. Many existing work such as SBGP [19] and SoBGP [25] relying on strong cryptography and PKI faces serious adoption difficulties. Several recent work [30, 10, 35] in this area attempt to reduce the computational overhead associated with these solutions, another obstacle to wide adoption. Compared to existing network-based, incrementally deployable mitigation schemes such as PG-BGP [18] and ACR [32], our mitigation scheme is complementary and identifies a more effective attack defense scheme that achieves the benefit close to global adoption with only partial deployment.

3. METHODOLOGY

In this paper, we study the proposed defense schemes us-

Table 1: Taxonomy of prefix hijacking defense techniques.

Defense		Network-based	End-host-based
Detection		MOAS [36], geo [20], PHAS [22], fingerprinting [16], hop-count [37], routing information objects [28]	ACR [32]
Reactive		Manual response to install filters, ACR [32], MIRO [33], <i>route purging-promotion</i>	Overlay routing, <i>e.g.</i> , RON [7]
Proactive	Crypto.-based	S-BGP [19], So-BGP [25], SPV [17], listen-whisper [29]	-
	Non-crypto.-based	PG-BGP [18], intentional deaggregation, bogon filter, Hi-BGP [27], <i>customer route filtering</i>	-

ing simulations on inferred AS topologies¹. Before presenting the defense schemes, we first discuss our methodology.

We obtained an AS topology annotated with AS relationships by running Gao’s algorithm [14] on BGP routing table dumps collected from around 70 vantage point ASes via RouteViews [2]. The topology contains 23,289 ASes, 55,352 inter-AS edges including 44,315 provider-customer (p2c) edges, 543 sibling-sibling (s2s) edges, and 10,494 peer-peer (p2p) edges. We also used a recent topology from CAIDA [13], and found that simulations on those two topologies produce similar results. For the rest of this paper, we present the results only on our inferred topology.

Our simulator emulates BGP route update propagation and the BGP decision process. The routing policies are configured at each AS based on AS relationships. Route selection policy is that firstly route selection is *profit-driven*, *i.e.*, customer’s route is preferred over peer’s route, and peer’s route is preferred over provider’s route; secondly a shorter route is preferred when there is a tie policy-wise; finally the AS number is used to break the tie length-wise. Route export complies with AS relationships. The same routing policy model has been used in previous studies [23, 18, 32].

We note that although a recent work [24] has proposed an AS topology model shown to predict AS paths with considerable accuracy, the model is not suitable for simulating prefix hijacks. The policies in this model are trained in the scenario where the victim originates the prefix, but not the scenario where attacker originates the prefix. In other words, the policies dictating the propagation of the attacker’s bogus routes are not captured by the trained policies. As a result, the model can not well predict the propagation of attacker’s bogus routes.

4. REACTIVE DEFENSES

As discussed earlier, prefix hijack detection is only the first step towards fully automated defense. Detection-based response today relies on human intervention, which is slow and error-prone. In this section, we propose a *reactive, detection-assisted mitigation scheme* that automatically responds to detected prefix hijacks and hence mitigates the adverse impact of the attacks in a timely fashion.

We make the following assumptions on the prefix hijack detection system used to assist automated hijack mitigation.

¹Both our topology and simulator are available at <http://www.ece.purdue.edu/~zhang97/prom/>

The fingerprinting-based detection system [16] and the RIPE MyASN service [3] meet all these requirements.

1. *Real-time detection*. The detection lag limits the benefit of mitigation.
2. *Low false-positives*. Mis-identified hijacks can degrade routing of relevant prefixes.
3. *Victim and bogus route identification*. This guides the our mitigation system to take effective mitigation response.

4.1 Mitigation System Overview

Our proposed mitigation system extends a prefix hijacking detection system with a set of counter-measure actions upon detecting a prefix hijack. It does so by contacting a set of preselected *lifesaver ASes* and instructs them to take one or two possible actions to revert the polluted routing tables in these ASes and in other ASes. The mitigation system is trusted by the lifesaver ASes, and receives a live BGP feed from each lifesaver AS to guide its decision. The trust between mitigation system and lifesaver ASes will be further discussed in Section 4.6.

Ideally, all ASes in the Internet participate and act as lifesaver ASes to completely eliminate the bogus routes; however, it is difficult to achieve such global adoption. In practice, the lifesaver ASes are typically large ISPs traversed by many network paths, which have more incentives for deploying security features. The mitigation actions executed by the lifesaver ASes remain effective until the original bogus route is withdrawn, at which point the mitigation system instructs the lifesaver ASes to revert to the previous state before the attack.

The mitigation system operates as follows. Upon detecting a prefix hijack, the detection system notifies the mitigation system about the hijack with three pieces of information: the attacker AS, the victim AS, and the victim prefix. Such information allows any AS (any routers) to differentiate between bogus routes which end with the attacker AS and valid routes which end with the victim AS. The mitigation system then contacts and instructs the lifesaver ASes to perform one or two possible actions described below:

- *Bogus route purging*. Each lifesaver AS deletes the bogus routes from its routing table. Given such ASes are typically large ISPs, the bogus route propagation is throttled. Similar to conventional manual response, bogus route purging blocks propagation of

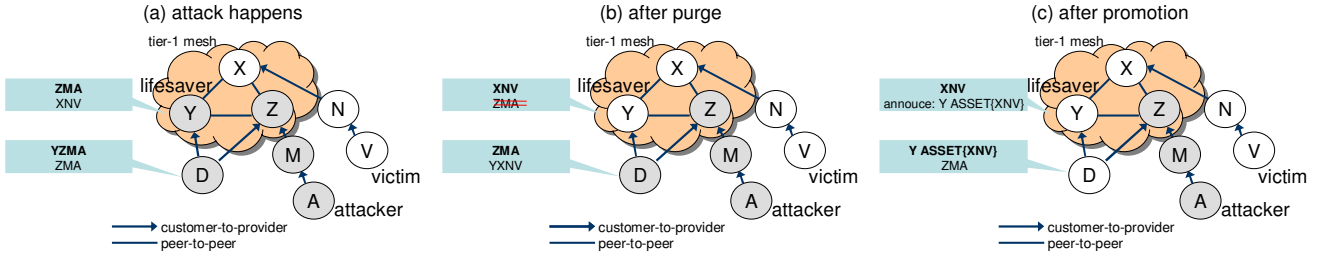


Figure 1: Example of purging-promotion. Gray nodes are polluted ASes. The boxes show the routing state of the ASes: including the routes learned, and the routes adopted (in bold), and the routes announced to neighbors.

bogus routes by deleting it. This is beneficial with even just a few well-connected ASes taking this action. However, ASes that still receive the bogus route may prefer the route over valid route based on BGP’s route selection decision process.

- *Valid route promotion.* A selected subset of lifesaver ASes are chosen by the mitigation system to further perform route promotion for the valid route to the victim AS: each selected promoter AS modifies the valid route by moving all ASes in the AS_PATH into an AS_SET. The AS_SET attribute is a mechanism used for route aggregation [5, 6] and effectively shortens the AS path to a prefix². By exploiting AS_SET, route promotion makes valid routes more attractive in the BGP best route selection process, since the AS path length is effectively shortened. To maximize the promotion effect, the promoter AS announces to all its neighbors the shortened promotion route, as if the victim prefix is the promoter’s own prefix.

Figure 1 shows an example of prefix hijack and how purging-promotion helps to mitigate the attack. Due to space limit, we scale a realistic scenario down to a small-size scenario consisting of three tier-1 ASes, one of them being the lifesaver, and several tier-2 and tier-3 ASes. In Figure 1(a), A hijacks V’s prefix, making Y, Z, M and D polluted by bogus routes. The lifesaver Y then attempts to revert the routing tables of the polluted ASes using purging and promotion. Y has learned both a valid route **ZMA** and a bogus route **XNV** and thus easily reverts itself by *purging* the bogus route **XNV** (Figure 1(b)). All of Y’s single-homed customers are reverted as well. Furthermore, Y can revert its multi-homed customer D by *promoting* Y’s route (Figure 1(c)), *i.e.*, put **XNV** into AS_SET construct to make this new route adopted by D, since the route appears to be shorter than before.

Path change due to promotion. Route promotion makes more ASes adopt valid routes. However, these valid routes may not be the same valid routes before the hijack occurred.

²BGP protocol specifies that AS_SET contributes only one to the path length no matter how many ASes are in AS_SET.

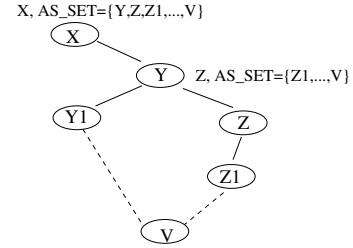


Figure 2: Prolonged path with multiple promoters. Both AS X and AS Z promote routes to V.

As a side effect of promotion, some AS that were not polluted may select valid routes that are actually longer but allegedly shorter than before the hijack. That is, these ASes experience prolonged paths. We show in Section 4.4 that this path inflation is not significant. We believe that eliminating bogus routes is of primary concern in mitigation prefix hijacking, because after all bogus routes cause unreachability while prolonged paths do not.

How many promoters? In selecting how many lifesavers to perform route promotion, there is an intricate tradeoff between reduced route pollution and the quality of reverted valid route. On one hand, using more promoters leads to reverting more polluted ASes to use promoted valid routes. On other hand, using multiple promoters can lead to prolonged valid route back to the victim, as shown in Figure 2. Assume before a hijack, AS Y uses a shorter path going back to an offspring customer V (victim AS) via Y1 than via Z. When both X and Z are performing route promotion for V, AS Y switches to advertising an allegedly shorter route **YZ...V**, which actually is longer than the original route back to V via Y1. Then AS X also switches to advertising an even longer but allegedly shorter route **XYZZ₁...V**.

Promoter selection. To qualify as a promoter AS, an AS needs to be pollution-free, either by itself, or by purging in case that it has at least one neighbor AS that is not polluted. This is because otherwise the promoter AS cannot forward data packets back to the victim, and hence has no valid route to promote. The promoter is selected by the mitigation system using the strategies discussed in Section 4.3 soon after the detection of attack and all lifesavers have per-

formed purging. However, if the hijack is detected by the detection system before the bogus route has converged in the global network, it is possible that after the promoter has been selected, the pollution further spreads out and pollutes all neighbors of the promoter, violating this qualification condition. In this case, the mitigation system will re-select another qualified promoter to replace. In our evaluation study in Section 4.4, we found that bogus route purging with lifesavers chosen using strategies discussed in Section 4.3 ensures most lifesavers to be pollution-free, making promoter re-selection unlikely to occur.

Protocol implication. We note that route promotion does not violate the BGP protocol, as it is a special route aggregation on the original route. Route promotion is the opposite of AS path prepending, a widely-used technique for making routes less preferred by prepending one’s AS to the AS path more than once. Both approaches attempt to influence route selection of other ASes by adjusting the AS path length without violating forwarding integrity of ensuring packets still reaching the correct destination.

Although promotion complies with BGP protocol, promoting a route causes temporary deviation from the AS relationship between the promoter and its neighbors. This means that promotion creates a new AS relationship other than the traditional customer-provider and peer-peer relationships. Therefore, we study the implication of promotion on route convergence guarantee and delay in Section 4.2.

Note that while our automated purging-promotion scheme provides timely mitigation against prefix hijacking, using a handful lifesavers does not always eliminate the bogus route across the entire Internet. In principle, the propagation of bogus routes can be blocked more effectively by choosing the handful lifesavers to be close to the attack router. However, this assumes that we have a large number of lifesaver candidates. Therefore, our proposed automated scheme is not a substitute for the traditional manual response whose goal is to remove the offender. Instead, our scheme complements the traditional manual response by quickly removing the impact of prefix hijacking from a large majority of the networks.

4.2 Correctness and Performance Analysis

We show the proposed route promotion scheme will preserve the convergence properties of the current BGP.

A previous study [15] proposed several guidelines of routing policies that guarantees Internet routing convergence without global coordination, and a later work [31] further showed that these guidelines are typically followed on the Internet. We add our route promotion to this framework, and prove the guaranteed convergence under the assumption that these guidelines are followed. Note that without this assumption, the basic routing convergence, *i.e.*, without route promotion, is not guaranteed on the Internet.

CLAIM 1. *For a BGP system that has only customer-provider and peer-to-peer relationships, and multiple route*

promoters, if all ASes follow the route preference guideline in [15], then the system is safe.

Our proof is an extension of the proof in [15], also based on the same two lemmas as Lemma 5.1 and Lemma 5.2 in [15].

LEMMA 1. *The BGP system has a stable state.*

PROOF. We construct an activation sequence σ^* that leads to stable state as follows.

To describe the sequence, we use the $S(i)$ to represent a linear ordering of ASes that starts with AS i and conforms to the partial order in customer-to-provider DAG, concatenated by another linear ordering of ASes that conforms to the partial order in provider-to-customer DAG (a combination of phase 1 and phase 2 in [15]).

In addition to normal activations, our constructed sequence also contains meta activations that set the modes of promoters. A promoter has two modes, *locked* and *unlocked*. In locked mode, a promoter acts as a normal AS, and in unlocked mode, a promoter performs promotion. Initially, all promoters are in locked mode. They remain locked until the meta activation $unlock(i)$ sets the mode of promoter i to unlocked.

Given a victim AS v and n promoters p_1, p_2, \dots, p_n , our constructed activation sequence σ^* consists of the following $(n + 1)$ phases:

phase 0: $S(v)$
 phase 1: $unlock\ p_1, S(p_1)$
 phase 2: $unlock\ p_2, S(p_2), S(p_1)$

 phase i : $unlock\ p_i, S(p_i), S(p_{i-1}), \dots, S(p_1)$

 phase n : $unlock\ p_n, S(p_n), S(p_{n-1}), \dots, S(p_1)$

It can be proved by induction that after phase i , promoters $1, 2, \dots, i$ are in unlocked mode, and the system reaches a stable state. (Details are be found in [34]). As a consequence, after phase n , all promoters are fully functioning because they are in all unlocked mode, and the system reaches a stable state. This proves Lemma 1. \square

LEMMA 2. *The BGP system converges to the stable state for any initial state and any fair activation sequence.*

Lemma 2 can be proved by extending the proof in [15].

CLAIM 2. *Let $MinRouteAdver\ period$ be Δ . The convergence time of a route promotion of a IP prefix by one or more ASes is at most $\Delta \cdot D$, where D is the longest simple path of ASes which is bounded by the number of ASes in the network. The number of route update messages generated during convergence is bounded by $(D \cdot E)$, where E is the number of BGP session between the routers.*

The convergence time for the single-promoter case is the same as in the unmodified BGP. The main reasoning for the convergence time with multiple promoters staying the same is based on the following observation, which is the same as Observation 2 in [21].

OBSERVATION 1. *The primary effect of a MinRouteAdver timer is to impose a monotonically increasing path metric for successive k -level iterations (convergence rounds).*

PROOF. We separate two cases. In case 1, after convergence, no promoter ends up in the AS Set of other promoters’ advertised route (for the victim prefix.) In other words, the AS Set used to reach the victim AS when each promoter started advertising the promoted route is not affected by other promoters during convergence. This case is no different from the legitimate multiple-origin ASes for a prefix scenario in unmodified BGP. Hence the convergence time of this case stays the same as later.

In case 2, after convergence, some promoters ends up in the AS Set of some other promoters’ advertised route (for example, Figure 2). We define a partial ordering of the promoters based on this relationship: if promoter p_i appears in the AS Set of promoter p_j , then $p_i < p_j$. One can then construct a forest of all the promoters using topological sort based on the partial ordering.

With this relationship, the overall convergence of multiple promoters advertising routes using AS Set can be reasoned as follows. We assume there is only one tree in the forest as multiple trees do not interference with each other (a simple generalization of case 1). First, the promoted route of the tree root is propagated, savaging all the ASes reached that preferred the new route. When the announcement reaches its child promoter(s) in the tree, the AS Set of the child promoter is updated, and the child promoter advertises the new shorter route for the victim prefix (because the path to the parent promoter is shorter than that to the victim AS.) This new advertisement should not affect any ancestor promoters or any ASes that have already switched to their final routes (routes to the victim prefix after global convergence.) The propagation process continues and eventually reaches the leaf promoters in the tree. Again, they update their routes for the victim prefix and advertise the updated routes. From now on, no promoters’ route will ever be affected, and hence the scenario is no different from the legitimate multiple-origin ASes scenario. Hence the total convergence time is at most $\Delta \cdot D$, where D is the longest simple path of ASes. \square

We note that the promoter ASes typically reside in the core of the Internet with only a few AS hops away from most other ASes. Thus the convergence delay for the promotion route, which is also the latency for the mitigation action to take effect, is expected to be quite low. An empirical study [26] shows that convergence delay of routes originated from the core is typically observed from most ASes as less than one minute. We expect that the mitigation latency of our scheme is typically within one minute.

4.3 Lifesaver and Promoter Selection Strategies

Since the route purging is performed on all lifesavers while route promotion is performed on one or a few lifesavers, the effectiveness of our mitigation scheme are deter-

mined by both the strategy of selecting lifesaver ASes among the ASes in the Internet when deploying the mitigation system and the strategy of selecting the promoter AS among these lifesaver ASes when a prefix hijack is detected.

The selection of lifesaver ASes affects the effectiveness of bogus route purging. The selection is challenging because they are selected prior to attacks whose locations are not yet known. Intuitively, choosing the lifesavers among the most well-connected ASes would best throttle the propagation of bogus routes and hence maximize the benefit.

The selection of promoter directly affects the effectiveness of valid route promotion. In valid route promotion, the promoter effectively “takes over” the victim prefix from the victim AS and announces the prefix as the promoter’s own. This behavior is analogous to the case where the promoter’s own prefix is hijacked by the attacker. So the benefit of valid route promotion is closely related to the promoter’s resilience against the attacker, *i.e.*, how well the promoter can protect its own prefix against the hijack. Therefore choosing the most resilient AS against the attacker maximizes the effectiveness of valid route promotion. Intuitively, well-connected tier-1 ASes have shorter paths to the other ASes, and hence are generally more resilient. However, a recent work [23] has shown using simulations that the most resilient ASes are tier-2 ASes with large numbers of providers mainly due to profit-driven routing policies on the Internet. Furthermore, because the selection of lifesavers dictates where the promoter comes from, resilience is also considered in the lifesaver selection strategy.

We propose several practical selection strategies as listed in Table 2 and Table 3. Lifesaver selection occurs during deployment, and is therefore based on static AS topological properties. One strategy is to use the node degree which indicates an AS’s connectivity. Another is based on the number of providers of a tier-2 AS which reflects that AS’s resilience. Promoter selection happens after attack detection, and hence uses information on the victim and the attacker. For example, the *near* strategy aims at preventing the neighborhood of the attacker from pollution and thus limiting the scope of the attack. The *far* strategy aims at maximizing the route length reduction from the original route to the promotion route. Finally, we include “optimal” which represents the best possible promoter selection strategy based on simulations. For this strategy study, we focus on selecting single promoter to gain some insight on its impact on the mitigation benefit, but we also include a simple strategy *all* that use all lifesavers as promoters.

In the following, we use the notation “ $xxx|yyy$ ” to denote the combined strategy, where xxx is the lifesaver selection strategy and yyy is the promoter selection strategy.

4.4 Evaluation

We evaluate our proposed scheme using simulations on the inferred AS topology (Section 3). N ASes on the AS topology are chosen as lifesavers using different strategies.

Table 2: Lifesaver selection strategies.

Name	Description
<i>degree</i>	Select the largest-degree ASes as lifesaver ASes.
<i>resilience</i>	Select tier-2 ASes with the largest number of providers as lifesaver ASes.
<i>hybrid</i>	Select half of lifesaver ASes using strategy <i>degree</i> , and select the other half using strategy <i>resilience</i> .

Table 3: Promoter selection strategies.

Name	Description
<i>random</i>	Randomly select a lifesaver as long as it has not been polluted.
<i>far</i>	Select the lifesaver that has not been polluted and is farthest from the victim in terms of AS path length.
<i>near</i>	Select the lifesaver that has not been polluted and is nearest to the attacker in terms of AS path length.
<i>tier2-rand</i>	Randomly select a promoter among the unpolluted tier-2 lifesavers if there is any. Otherwise, randomly select among all unpolluted lifesavers.
<i>optimal</i>	Select the lifesaver whose promotion action achieves the largest pollution reduction.
<i>all</i>	Select all lifesavers as promoters. This is used to defend against colluding attacks.

We vary N from 0 to 24. For each N , 200 random regular prefix hijack trials are simulated. For each trial, a single attacker AS and a single victim AS are randomly selected among all Internet ASes. Stealthy hijacks using falsified AS paths are not considered, because they complicate hijack detection but not mitigation. Handling subprefix hijacks is discussed later in Section 4.6.

4.4.1 The Benefit of Bogus Route Purging

We first study the benefit of bogus route purging alone. Figure 3(1) shows the benefit of bogus route purging with various numbers of lifesavers chosen by the three strategies in Table 2. The figure shows that purging bogus routes at a few ASes provides some protection against prefix hijack. This is because of the route diversity at these lifesavers. The well-connected lifesaver has many neighbors that provide diverse routes to a destination prefix. It is unlikely that all these neighbors are polluted, and hence the lifesaver is highly likely to find a valid route. Also note that the *degree* strategy performs better than the other two strategies which tend to choose ASes with smaller degree. Therefore, maximizing the degree of lifesavers achieves the best bogus route purging benefit.

4.4.2 The Benefit of Route Purging-Promotion

Next we study the benefit of combining bogus route purging and valid route promotion. We assume the *degree* strategy as lifesaver selection strategy. We assume a single route promoter, and study the first four promoter selection strategies in Table 3, excluding the *all* strategy to isolate the effects of multiple promoters from the impact of selection strategies. Figure 3(2) shows the benefit of route purging-promotion using these strategies as well as using purging alone. We make the following observations. First, route

purging-promotion achieves higher benefit than bogus route purging alone. In Figure 3(2), with 8 lifesaver ASes, the fraction of Internet ASes that are polluted by a hijack is reduced from 50% to 20% using promoter selection strategy *random*, whereas this fraction is 30% for purging alone (Figure 3(1)). Second, in Figure 3(2), there is a gap between those three strategies and *optimal*. The gap is because path length is not the only deciding factor in BGP decision process. Local preference dictated by AS relationship overrides path length. Length-based strategies *Far* and *near* as well as *random* do not effectively capture the resilience of the optimal promoter ASes. Actually, we found that the optimal promoters are mostly tier-2 ASes. This observation motivates using the resilience-aware strategies listed in Table 2 and Table 3.

4.4.3 Enhancement by Resilience-based Strategies

Next we evaluate the effectiveness of several combined lifesaver and promoter selection strategies, again assuming a single promoter. Our evaluation includes four combined strategies: *degree|random*, and three resilience-aware strategies, namely *degree|tier2-random*, *resilience|random*, *hybrid|tier2-random*. The results are shown in Figure 3(3). We make the following observations. First, *resilience|random* performs worst. Although choosing lifesavers based on resilience maximizes the benefit of valid route promotion, this benefit is offset by the inferior benefit of bogus route purging by these lifesavers. It has been shown in Figure 3(1) that maximizing the degree of the lifesaver ASes achieves the most effective bogus route purging. Second, *degree|tier2-random* and *hybrid|tier2-random* perform best. They both trade off between maximizing connectivity for purging and maximizing resilience for promotion.

4.4.4 Prolonged Routing Paths due to Promotion

A negative effect of route promotion is potentially sub-optimal route selection. The route promoter can oversell its route, *i.e.*, when the actual length of the promotion route is longer than the length calculated in BGP decision process. Figure 4(1) shows the AS path inflation experienced by the pollution-free ASes in route purging-promotion using the *degree|random* strategy and a single promoter. The *path inflation* is defined as the relative AS path length increase experienced by each AS after the promotion compared to the original AS path length. We observe that the AS path inflation is mostly small. In most cases more than 50% of the pollution-free ASes experience no inflation at all, more than 70% of the pollution-free ASes experience less than 20% inflation, and almost all pollution-free ASes experience less than 50% inflation.

We also analyze the tradeoff between reduced pollution and increased path inflation with more promoters as discussed in Section 4.1. We make every lifesaver perform both route purging and promotion, and vary the number of

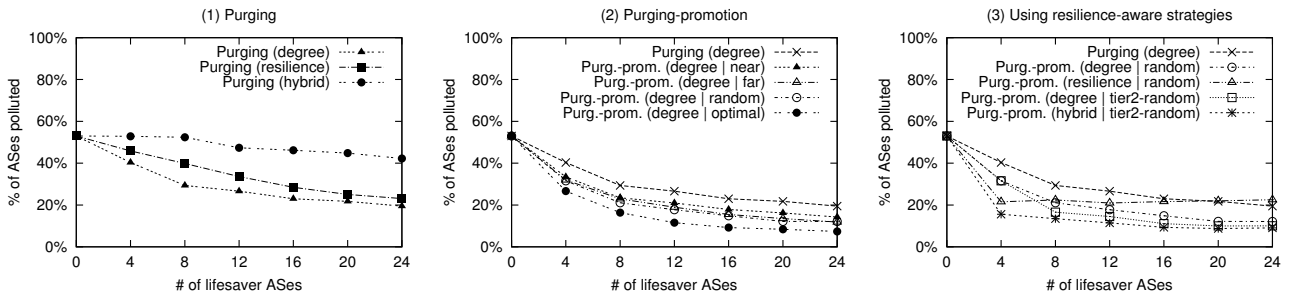


Figure 3: Pollution of a random prefix hijack when a number of lifesavers perform (1) bogus route purging, (2) purging-promotion using degree-based strategies, (3) purging-promotion using resilience-aware strategies.

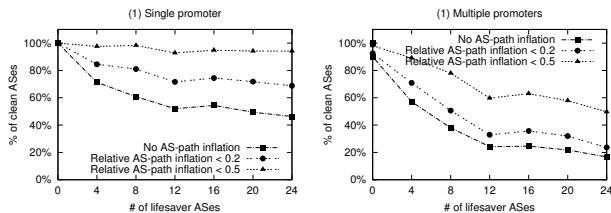


Figure 4: AS-path inflation experienced by pollution-free ASes under purging-promotion. The strategy of selecting lifesaver and promotor is (1) *degree|random*, (2) *degree|all*.

lifesavers. Figure 5(3) and Figure 4(2) show that for a single attacker, more lifesavers results in fewer polluted ASes, but the path inflation for pollution-free ASes also increases drastically. Based on this tradeoff, a single promotor appears sufficient assuming the presence of one attacker.

4.4.5 Colluding Attack and Defense

So far we have assumed that the attacker originates a bogus route from a single AS. With access to multiple ASes, the attacker can maximize the adoption of bogus routes by originating a bogus route from each of these ASes. We now study the pollution of colluding attacks and how our mitigation system defends against these attacks. We vary the number of attacker ASes from 1 to 5.

Figure 5(1) shows the pollution of such colluding attacks when all lifesavers perform purging. Purging is less effective against colluding attacks than regular attacks. An interesting observation is that the lifesavers often lose the combat against attacker ASes even when the lifesavers outnumber the attackers. For example, a 4-AS colluding attack pollutes more than 50% of the Internet even with 8 lifesavers. This is because the “machinery” used by two sides are different. Attacker ASes originate routes, while lifesaver ASes delete routes, which is far less effective.

Figure 5(2) shows the pollution when a single lifesaver performs promotion in addition to purging. This is the strategy shown to be effective to handle a single attacker AS. However, with multiple attackers, the pollution reduction is small compared to the corresponding purging-alone cases.

To more effectively promote valid routes in the presence of multiple attackers, we have all the lifesavers perform promotion. Each lifesaver does promotion independently, and thus no global coordination is needed. Figure 5(3) shows the dramatic improvement in the defense effectiveness. Given that colluding attacks have never been witnessed on the Internet, selecting single promotor is sufficient currently, because of its simplicity, fast convergence, and minimal sub-optimal routing.

4.5 Implementation

The mitigation system can be implemented in software, very similar to the setup of the Routing Control Platform (RCP) [11] which is used to control the route selection decision of routers within a single ISP. The mitigation system communicates with the RCP-like system in each lifesaver AS to instruct the AS to perform route purging and promotion. If the lifesaver AS does not deploy centralized route management using a system like RCP, the mitigation system needs to directly communicate with one router in each lifesaver AS. That router in turn distributes the updated routing information to other routers inside the AS relying intradomain routing hierarchy such as iBGP mesh and route reflector based structure.

4.6 Discussions

In the following, we discuss some issues including trust, detection accuracy, and deployment.

Detection accuracy. The effectiveness of our reactive schemes rely on accurate hijacking detection, which is a research topic of its own, and expects improvement in the future though good progress has been made recently. However, our scheme does not require perfect detection accuracy. False positives just cause the traffic to take a slightly longer path than before. False negatives are of more concern. We will investigate them in our future work.

Workload on lifesaver ASes. Performing route promotion increases the workload of the lifesavers, but we expect this increase to be negligible. Given that prefix hijacking is a rare event, the number of prefixes under attack simultaneously is small compared to the large number of prefixes that these large lifesaver ISPs provide transit for. Therefore, the extra

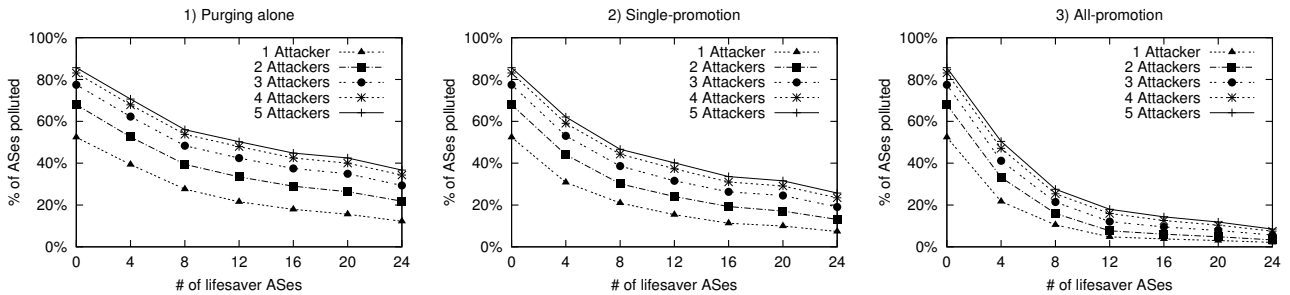


Figure 5: Pollution of a random colluding hijack when (1) all lifesavers perform purging, (2) single lifesaver performs promotion in addition to purging, (3) all lifesavers perform promotion in addition to purging.

workload for helping hijacked prefixes should be negligible.

Deployment incentive. Route purging-promotion provides strong incentive for deployment, because a lifesaver provider prevents its customer from being hijacked and being polluted by hijacks. Moreover, the scheme is effective with small scale of deployment.

Trust between mitigation system and lifesaver ASes. We assume that there will be one well-known detection and mitigation system, which does not misbehave, *i.e.*, it is a public Internet service like the DNSBL used for blacklisting spamming hosts. Lifesaver ASes and the mitigation system authenticate each other using SSL certificates when the mitigation system collects valid BGP feed or sends purging-promotion instructions. There are a limited number of entities whose identities need to be verified. Each lifesaver needs to verify the identity of the mitigation system, but not other lifesavers. The mitigation system needs to verify the identities of the lifesavers. Hence our scheme does not require a full-blown PKI like what is required in SoBGP [25].

Subprefix hijacking. Route purging-promotion could be extended to handle subprefix hijacks. Upon the detection of subprefix hijacks, the detection system notifies the victim AS. If the victim AS could originate the hijacked subprefix promptly, the subprefix hijacks is no different from a regular prefix hijack.

4.7 Summary

We have presented a reactive mitigation system combining bogus route purging and valid route promotion. Simulations show:

- Purging bogus routes at a few high-degree ASes (*e.g.*, 20 highest-degree ASes) provides good protection against prefix hijack (*e.g.*, a reduction of pollution down to 24%). Maximizing the degree of lifesavers achieves the best bogus route purging benefit.
- Adding promotion to purging reduces the remaining pollution by 33% ~ 57%.
- Selecting lifesavers and promoters by trading off between maximizing connectivity for purging and maximizing resilience for promotion achieves the best benefit.

- The resulting routing sub-optimality is insignificant. More than 50% of the pollution-free ASes use AS paths of the same length, and almost all of them adopt AS paths less than 50% longer compared to before the attack.

5. PROACTIVE DEFENSES

The reactive mitigation scheme proposed in Section 4 relies on an accurate hijack detection system, as it is triggered after a hijack is detected. However, the detection system may not detect all attacks due to the limited visibility. In this section, we first study the coverage of the detection system to motivate the need for proactive prevention schemes. We then analyze the effectiveness of a known proactive scheme: customer filtering.

5.1 Detection Evasion

We define attack detection evasion as follows.

DEFINITION 1. (Detection Evasion) We denote the monitoring system as $SM = m_1, m_2, \dots, m_n$, where there are altogether n monitors in distinct ASes. Given an attacker A , a victim V , and the hijacked prefix p , if $\forall i, Pref_{m_i}^A(p) < Pref_{m_i}^V(p)$, where $Pref_{m_i}^A(p)$ is the route preference value for p announced from A observed by monitor m_i , then attacker A can hijack V 's p without being detected.

Note that since the detection system receives the best route from each monitor, only when at least one of the monitors chooses the bad route as its best route, hijacking becomes visible to the monitor system.

An example of attack evasion from the monitoring system is depicted in Figure 6. Attacker A hijacks one of victim V 's prefix p . Node M is the monitoring system. We present the system as a single node for ease of explanation. M receives both routes for prefix p originated from A and V with different path length. Obviously, due to route selection based on the commonly used profit-driven policy, *i.e.*, preferring customer over peer and over provider, M selects the route from V due to preference for customer routes.

We summarize the conditions for attack evasion.

OBSERVATION 2. An attacker can evade detection if any of the following is true for all monitoring nodes.

- The victim route is a customer or peer route, and the attacker route is a provider route.
- The victim route is a customer route, the attacker route is a provider or peering route.
- Both the victim and attacker routes have the same profit-driven preference, but the victim route is shorter.

We further perform simulations to demonstrate the real evasion threat under the RouteViews monitoring system, commonly used by many studies. For scalability, we ignore stub AS nodes which do not provide transit in the simulations. Our results can be easily extended to consider stub nodes which have to traverse through one of their providers to reach other networks. We identified 27,145 attacker/victim pairs evading detection, accounting for 0.2% of all possible AS pairs ignoring stub nodes. Among them there are 2194 distinct attackers and 1691 distinct victims. Figure 7 further shows how many possible victims a given attacker can choose to hijack, and similarly how many possible attackers can affect each victim. Among these potential attackers, 72% are edge ASes (tier-4, tier-5). Similarly, 73% of the victims are edge ASes.

Although an attacker can evade detection by carefully selecting victims, this limits attack flexibility. There is a clear trade-off between the ability to pollute many different ASes and the desire to evade detection. Figure 8 shows the fraction of polluted ASes from all evasion scenarios studied. We observe that 40% attacker ASes can only pollute 10% of all the ASes if evading detection.

5.2 Customer Route Filtering

Section 5.1 shows that a hijack detection system relying on BGP feeds due to limited visibility cannot detect all possible prefix hijacks as needed by reactive schemes such as route purging-promotion. In this section, we study customer route filtering, a known proactive scheme that does not rely on real-time IP prefix hijack detection.

5.2.1 Design

Customer route filtering is currently practiced by several large ISPs to prevent their customers from injecting bogus routes. Such an ISP AS P maintains a route registry among P and its direct customers P_i . Each P_i registers the prefixes that P_i announces to P . These prefixes are prefixes originated by ASes in P_i 's customer-cone, i.e., by P_i , P_i 's customers, P_i 's customers' customers, and so on. Route filtering is performed at the each BGP session between P and its direct customer P_i . Any route announced by P_i for a prefix not registered is blocked by the filter at P .

While route filtering is potentially effective, ISPs performing route filtering rely on up-to-date route registries. In practice, ISPs can maintain route registries separately, e.g., as used by Level 3, or share one route registry, e.g., as used in RIPE [1]. In either case, the freshness of route registries is critical to route filtering. Although the local registry is relatively easier to maintain than a global registry as the partici-

pants are involved in direct business relation, it still requires coordinated efforts to synchronize the registries owned by different providers. Reassignment or new assignment of a prefix requires updates to the multiple registries of all the providers of the prefix owner that are higher up in the AS hierarchy.

5.2.2 Evaluation

Although customer route filtering has been practiced by some large ISPs, its effectiveness in defending against prefix hijacks has not been studied before, especially for partial deployment. Furthermore, it is unlikely to be voluntarily deployed globally, as it requires additional management overhead of keeping track of addresses allocated to customers whose multihoming practice further complicates it. In the following, we evaluate the effectiveness of partially deployed customer route filtering over the Internet. As in the previous experiments in Section 4.4, we randomly choose attacker and victim ASes, and simulate regular prefix hijacks. We consider the same degree heuristic used for route purging: the ASes with the largest degree are selected first. Selecting the ASes based on their resilience is not considered as the selected ASes do not originate new routes.

The solid triangle curve in Figure 9 shows the pollution by random prefix hijacks under customer route filtering. We see that customer route filtering provides limited protection against prefix hijacks. With 16 most well-connected ASes performing filtering, the fraction of polluted Internet is 32%, and with 80 performing filtering, the fraction is reduced to 9%. However, these numbers are much worse compared to route purging, with the same numbers of participating ASes due to these two reasons.

1. Customer route filtering is performed on limited links, whereas route purging are performed at the AS level. In the former case, an AS does not perform filtering on links to its peers or providers, it may import a bogus route. In contrast, an AS that implements route purging never imports bogus routes.
2. Even links that perform filtering cannot distinguish certain bogus routes: if both the attacker and victim are within the same customer-cone of the customer end of a link that implements the filtering, the filter is not effective. Such a link is considered to be *defense-incapable* for these attacks.

Figure 10 quantifies how often the above case (2) occurs. We define *defense-responsible* c2p link in a prefix hijack as a c2p link that satisfies the following two conditions: (1) the provider end of this link performs filtering; (2) this link is traversed by a normal route originated by the attacker. In other words, defense-responsible c2p links are those links responsible for defending against the bogus routes. We see that although the number of defense-responsible c2p links are seemingly large, 80% are defense-incapable.

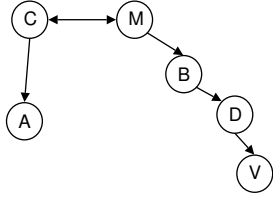


Figure 6: An example attack that evades detection.

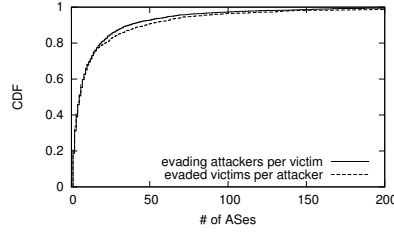


Figure 7: The number of attackers and victims under detection evasion.

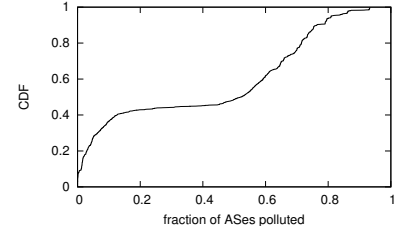


Figure 8: The polluted ASes under detection evasion.

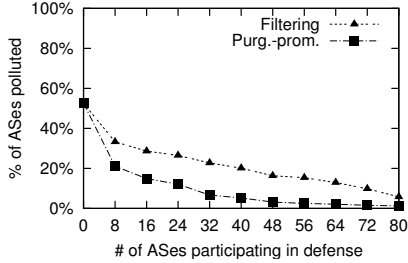


Figure 9: The pollution by hijacks with partially deployed with router filters and purging-promotion.

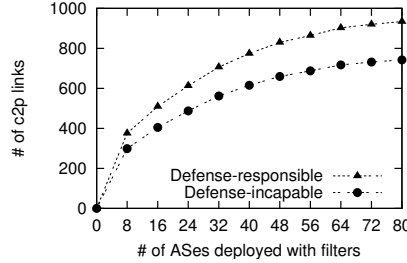


Figure 10: The capability of defensive c2p links in customer route filtering.

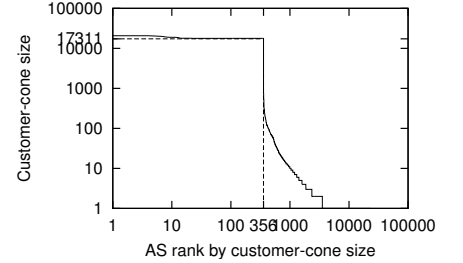


Figure 11: Customer-cone size of ASes.

The vast majority of defense-incapable c2p links is explained by Figure 11 which shows the customer-cone sizes of all ASes. Probably due to the wide use of multi-homing, 356 ASes (denoted by set W) have a customer-cone size larger than 17000 and the remaining ASes (denoted by set W^C) generally have much smaller customer-cone size (less than 100). Consider the filter between a provider P and one of its direct customers P_i . If P_i is in W , it is likely that both the attacker and the victim are within the customer-cone of P_i , making the filter defense-incapable. If P_i is in W^C , it is likely that the attacker is not within the customer-cone of P_i , making the filter not defense-responsible.

However, the high percentage of defense-incapable c2p links is not a completely negative observation. Figure 10 shows that the number of defense-capable c2p links consistently increases with the number of ASes deployed with filters, which contributes to the decrease of hijack pollution in Figure 9.

Route purging-promotion and customer route filtering complement each other. The solid square-and-circle curves in Figure 9 show the effectiveness of using customer route filtering together with route purging-promotion deployed on four highest-degree ASes and together with route purging-promotion deployed on eight highest-degree ASes, respectively. They both show an additional reduction of pollution to the case of using customer route filtering alone.

5.2.3 Summary

We have evaluated customer route filtering, a proactive scheme currently practiced by some large ISPs. Our simulations show that the effectiveness of customer filtering against prefix hijacking is much lower than route purging with the

same scale of deployment. This observation is because a significant proportion of the filters are unable to confine the bogus routes originated from the customer-cone, which is caused by the rich connectivity of the Internet topology.

6. RELATED WORK

Existing work in the area of proactively defending against routing attacks mainly focuses on using strong cryptography or incremental solutions such as intentional deaggregation to proactively prevent against routing attacks as shown in Table 1. We note that besides deployment difficulties partly due to computational overhead and PKI requirement, solutions such as SBGP [19] and SoBGP [25] do not completely eliminate routing attacks such as IP prefix hijacking, as they authenticate the routing information and the origin of the route, but do not ensure the correctness of the entire AS path.

Our study focuses on incrementally deployable network-based solutions. Several existing solutions fall in this category, but all with serious limitations. For example, intentional route deaggregation refers to the practice of ISPs advertise many small prefixes within its address block for fear of subprefix hijacks. Such practice increases the already large routing table sizes and also do not guarantee valid routes will be preferred over bogus routes. A recent proposal of pretty good BGP [18] merely delays the selection of suspicious routes and as a side-effect increases the time to adopt legitimate new routes. Note that our study has so far focused on hijacking of allocated and advertised IP prefixes, as they cause more damage compared to hijacking of unallocated or bogon routes. Bogon filters [4] is an effective approach to avoid propagating such invalid routes. However, similar to ingress and customer route filtering, such filters

are not globally deployed.

Our work proposes automated reactive mitigation response through route purging and promotion, which is complementary to the current manual response to detected routing hijacks. Finally, our reactive mitigation system relies on an accurate and timely detection system, achieved from several existing systems [16, 20, 22, 37]. Our work is also motivated by a recent study [23] analyzing the resilience of Internet topology against prefix hijacks.

7. CONCLUSIONS

In this study, we address the defense against an important attack targeted at the current Internet routing system, namely the IP prefix hijacking attack against BGP, by developing novel incrementally deployable network-based reactive mitigation solution. Using our proposed solution, simulation results based on realistic network topologies demonstrate that with intelligent selection of deployment locations, the number of polluted ASes can be reduced down to around 15% with a relatively small number of participating ASes (e.g., 20). In contrast, the current network-based solution such as customer route filtering is much less effective at limiting the impact of polluted routes. We believe that our work explored the limits of readily deployable network-based defense against IP hijacking. We are also the first to point out the general limitations of hijack detection systems due to their reliance on BGP feeds and caused by evasion. These lessons illustrated by our work provide guidance for designing the secure next-generation Internet routing system.

Acknowledgments

We thank the anonymous reviewers and Olivier Bonaventure for their helpful comments. This work was supported in part by NSF grants CyberTrust-0430204, CAREER-0238379, and CAREER-0643612.

8. REFERENCES

- [1] RIPE Network Coordination Centre. <http://www.ripe.net/>.
- [2] Route Views Project. <http://www.routeviews.org/>.
- [3] The RIPE NCC MyASN service. <http://www.ris.ripe.net/myasn.html>.
- [4] The Tem Cymru Bogon Route Server Project. <http://www.cymru.com/BGP/bogon-rs.html>.
- [5] A Border Gateway Protocol 4 (BGP-4), Jan. 2006. RFC 4271.
- [6] BGP-4 Implementation Report, Jan. 2006. RFC 4276.
- [7] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris. Resilient overlay networks. In *Proc. ACM SOSP*, 2001.
- [8] H. Ballani, P. Francis, and X. Zhang. A Study of Prefix Hijacking and Interception in the Internet. In *Proc. ACM SIGCOMM*, 2007.
- [9] V. J. Bono. 7007 Explanation and Apology. NANOG email on Apr 26, 1997.
- [10] K. Butler, P. McDaniel, and W. Aiello. Optimizing bgp security by exploiting path stability. In *Proc. Computer and Communications Security (CCS)*, 2006.
- [11] M. Caesar, D. Caldwell, N. Feamster, J. Rexford, A. Shaikh, and J. van der Merwe. Design and Implementation of a Routing Control Platform. In *Proc. NSDI*, 2005.
- [12] H. Chan, D. Dash, A. Perrig, and H. Zhang. Modeling Adoptability of Secure BGP Protocol. In *Proc. ACM SIGCOMM*, 2006.
- [13] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, kc claffy, and G. Riley. AS Relationships: Inference and Validation. *ACM SIGCOMM CCR*, 37(1):29–40, Jan. 2007.
- [14] L. Gao. On inferring autonomous system relationships in the internet. *IEEE/ACM Trans. on Networking (TON)*, 9(6):733 – 745, Dec. 2001.
- [15] L. Gao and J. Rexford. Stable internet routing without global coordination. In *Proc. ACM SIGMETRICS*, 2000.
- [16] X. Hu and Z. M. Mao. Accurate Real-time Identification of IP Prefix Hijacking. In *Proc. IEEE Security and Privacy (Oakland)*, 2007.
- [17] Y.-C. Hu, A. Perrig, and M. Sirbu. SPV: A Secure Path Vector Scheme for Securing BGP. In *Proc. ACM SIGCOMM*, 2004.
- [18] J. Karlin, J. Karlin, S. Forrest, and J. Rexford. Pretty Good BGP: Improving BGP by Cautiously Adopting Routes. In *Proc. IEEE ICNP*, 2006.
- [19] S. Kent, C. Lynn, and K. Seo. Secure Border Gateway Protocol (Secure-BGP). *IEEE Journal on Selected Areas in Communications (JSAC)*, 18(4):582–592, Apr. 2000.
- [20] C. Kruegel, D. Mutz, W. Robertson, and F. Valeur. Topology-Based Detection of Anomalous BGP Messages. In *Symposium on Recent Advances in Intrusion Detection (RAID)*, 2003.
- [21] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian. Delayed internet routing convergence. In *Proc. ACM SIGCOMM*, 2000.
- [22] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang. PHAS: A Prefix Hijack Alert System. In *Proc. USENIX Security Symposium (Security)*, 2006.
- [23] M. Lad, R. Oliveira, B. Zhang, and L. Zhang. Understanding resiliency of internet topology against prefix hijack attacks. In *Proc. IEEE/IFIP Intl. Conf. on Dependable Systems and Networks (DSN)*, 2007.
- [24] W. Mhlbauer, A. Feldmann, O. Maennel, M. Roughan, and S. Uhlig. Building an AS-topology model that captures route diversity. In *Proc. ACM SIGCOMM*, 2006.
- [25] J. Ng. Extensions to BGP to Support Secure Origin BGP (soBGP), Oct. 2002. Internet Draft draft-ng-sobgp-bgp-extensions-00.
- [26] R. Oliveira, B. Zhang, D. Pei, R. Izhak-Ratzin, and L. Zhang. Quantifying Path Exploration in the Internet . In *Proc. ACM SIGCOMM IMC*, 2006.
- [27] J. Qiu and L. Gao. Hi-BGP: A Lightweight Hijack-proof Inter-domain Routing Protocol . Technical report, Univ. of Massachusetts , 2006.
- [28] J. Qiu, L. G. S. Ranjan, and A. Nucci. Detecting Bogus BGP Route Information: Going Beyond Prefix Hijacking . In *Proc. SecureComm*, 2007.
- [29] L. Subramanian, V. Roth, I. Stoica, S. Shenker, and R. H. Katz. Listen and Whisper: Security Mechanisms for BGP. In *Proc. NSDI*, 2004.
- [30] T. Wan, E. Kranakis, and P. van Oorschot. Pretty Secure BGP (psBGP). In *Proc. Network and Distributed System Security Symposium (NDSS)*, 2005.
- [31] F. Wang and L. Gao. Inferring and Characterizing Internet Routing Policies . In *Proc. ACM SIGCOMM IMC*, 2003.
- [32] D. Wendlandt, I. Avramopoulos, D. Andersen, and J. Rexford. Don't Secure Routing Protocols, Secure Data Delivery. In *Proc. ACM HotNets*, 2006.
- [33] W. Xu and J. Rexford. MIRO: multi-path interdomain routing. In *Proc. ACM SIGCOMM*, 2006.
- [34] Z. Zhang, Y. Zhang, Y. C. Hu, and Z. M. Mao. Practical Defenses Against BGP Prefix Hijacking. Technical report, Purdue University, 2007. <http://www.ece.purdue.edu/~zhang97/pub/prom.pdf>.
- [35] M. Zhao, S. W. Smith, and D. M. Nicol. Aggregated path authentication for efficient bgp security. In *Proc. Computer and Communications Security (CCS)*, 2005.
- [36] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. Wu, and L. Zhang. An Analysis of BGP Multiple Origin AS (MOAS) Conflicts. In *Proc. ACM SIGCOMM Internet Measurement Workshop*, 2001.
- [37] 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*, 2007.