Region-based BGP Announcement Filtering for Improved BGP Security

Fernando Sanchez
Florida State University
sanchez@cs.fsu.edu

Zhenhai Duan
Florida State University
duan@cs.fsu.edu

ABSTRACT

BGP prefix hijacking is a serious security threat on the Internet. In this paper we propose a region-based BGP announcement filtering scheme (RBF) to improve the BGP security. In contrast to existing solutions that indifferently prevent or detect prefix hijacking attacks, RBF enables differentiated AS and prefix filtering treatment and blends prefix hijacking prevention with deterrence. RBF is a light-weight BGP security scheme that provides strong incremental deployment incentive and better prefix hijacking deterrence. Experimental studies based on real Internet numbers allocation information and BGP traces show that RBF is a feasible and effective scheme in improving BGP security. For example, on the days without known BGP prefix hijacking attacks, only a small number of BGP announcements will be flagged as attacks. Importantly, by applying RBF to known BGP prefix hijacking attacks, we show that RBF can detect and filter both large-scale and small-scale BGP prefix hijacking attacks even if only a single prefix is hijacked.

Categories and Subject Descriptors

C.2.0 [**Computer-Communication Networks**]: General—*Security and protection*; C.2.3 [**Network Operations**]: Network monitoring

General Terms

Security, Reliability, Measurement

Keywords

BGP, BGP Security, Network Prefix Hijacking

1. INTRODUCTION

The Internet consists of tens of thousands of network domains or Autonomous Systems (ASes), each of which is a logical collection of networks under a common administrative control [9]. ASes exchange the reachability information of network prefixes via an inter-domain routing protocol. The current inter-domain routing protocol is the Border Gateway Protocol (BGP) [25, 28]. Despite the critical importance of BGP to the healthy operation of the global Internet system, BGP is inherently lack of security measures and is vulnerable to a number of security attacks [19, 20], including the highly publicized prefix hijacking attacks [3, 22, 26].

BGP prefix hijacking attacks take two broad forms: origin spoofing and path spoofing. In origin spoofing, an AS u originates a prefix (or a more specific prefix) belonging to another AS v without proper authorization from v. In path spoofing, an AS u announces a false route pretending to be en route to a prefix of AS v. In this case, the originating AS is still the prefix owner v. In this paper we focus on addressing the first type of prefix hijacking, i.e., origin spoofing. (We briefly discuss how the scheme developed in this paper can be extended to help address path-spoofing prefix hijacking in Section 5.) Prefix hijacking may occur because of either unintentional misconfiguration or intentional malicious attacks. To simplify the description we refer to both as *attacks*.

Given the importance of addressing the prefix hijacking problem to the global Internet security, many countermeasures have been developed to improve the BGP security, including both cryptographic [11, 16, 29, 30] and non-cryptographic methods [8, 10, 14, 17, 18, 23, 27, 31]. However, none of them have been widely deployed or adopted on the Internet, due to various reasons including high cost, lack of (partial) deployment incentives, and imprecise operational objectives, among others. Overall, they all try to indifferently prevent or detect prefix hijacking attacks, viewing all ASes and IP address prefixes equally. This indifferent view makes it challenging to incrementally deploy any BGP security schemes on the Internet [5].

In this paper we take a fundamentally different approach that provides differentiated AS and prefix treatment, and that blends prefix hijacking prevention with deterrence. We refer to this approach as region-based BGP announcement filtering (RBF). In RBF, the Internet is partitioned into regions; an AS in a region cannot

ASIACCS'10 April 13–16, 2010, Beijing, China.

originate a prefix allocated to a different region. Regions can be defined at different granularity depending on the security requirements of the ASes where RBF is deployed. In the current paper, we consider two region granularities: country-level and Regional Internet Registry (RIR) level [12]. The motivation behind RBF is the following. ASes within the same region normally have better communication channels, and more importantly, they are more likely to be within the same jurisdiction of law enforcement. By preventing prefix hijacking across different regions, we can confine all potential prefix hijacking attacks to the same region, where ASes can quickly resolve prefix hijacking caused by misconfiguration due to better communication channels, and more importantly, potential malicious attacks can also be largely deterred due to the explicit legal consequences involved in hijacking prefixes belonging to ASes in the same region.

RBF is a light-weight prefix hijacking prevention and deterrence scheme. It provides strong deployment incentives and can be incrementally deployed by individual ASes or regions [5]. RBF also has a precise operational objective of preventing cross-region prefix hijacking attacks to provide better attack deterrence. In this paper we present the design of RBF. We also evaluate the performance of RBF using the currently available AS and prefix *allocation* (in contrast to the voluntary registration) information maintained by RIRs. Our studies show that, on the days without known BGP prefix hijacking attacks, RBF will on average only flag 215 AS and prefix pairs as attacks each day, with the help of a small table to maintain the legacy cross-region BGP announcements. That is, network operators only need to examine a small number of potential hijacking attacks each day. Importantly, using BGP trace with known prefix hijacking attacks, we also show that RBF can detect and filter both large-scale BGP prefix hijacking attacks (where a large number of prefixes are hijacked) and small-scale attacks (where only a few or a single prefix is hijacked), should the attacks involve ASes and prefixes in different regions. Therefore, RBF provides ASes and regions with the confidence of what BGP prefix hijacking attacks may occur and enables them to respond to the attacks accordingly.

The remainder of the paper is organized as follows. In Section 2 we provide the necessary background on inter-domain routing and numbers allocation and assignment on the Internet. We present the design of the RBF architecture in Section 3. We perform experimental studies using real BGP traces in Section 4. Section 5 briefly discusses how RBF can be extended to handle path spoofing attacks, and summarizes related work. We conclude the paper in Section 6.

2. BACKGROUND

In this section we present the necessary background on BGP that is most relevant to our work. We refer interested readers to [25, 28] for a more complete description of BGP. In addition, we also describe the Internet numbers allocation and assignment including AS numbers and IP addresses.

2.1 Border Gateway Protocol

The Internet is a collection of tens of thousands of network domains or Autonomous Systems (ASes). Each AS has a unique AS number (ASN) and owns one or multiple IP address prefixes. ASes exchange network prefix reachability information using an inter-domain routing protocol, the Border Gateway Protocol (BGP). BGP has two types of route update messages—*announcements* or *withdrawals*. A route withdrawal, containing a list of network prefixes, indicates that the sender of the withdrawal message can no longer reach the prefixes. In contrast, a route announcement indicates that the sender knows of a path to a network prefix.

A BGP announcement message contains a list of *route attributes* associated with the destination network prefix. One important route attribute is *as_path*, the path vector attribute that is the sequence of ASes (i.e., their ASNs) that this route has been propagated over. We will use $r.as\_path$ to denote the *as_path* attribute of route r. Let $r.as\_path = \langle v_k v_{k-1} \dots v_1 v_0 \rangle$. The route was originated (first announced) by AS v_0, which owns the destination network prefix. Before arriving at AS v_k, the route was carried over ASes $v_1, v_2, \dots, v_{k-1}$ in that order. In a small portion of BGP routes on the Internet, an AS path *as_path* contains a set of ASNs *as_set*, in which ASNs do not have the particular order as discussed above. This is commonly caused by prefix aggregation. We discuss how to handle *as_set* in the next section.

After learning a set of candidate routes from neighbors, AS v selects a single *best* route to reach the destination, based on some local route selection policy [6]. AS v then propagates the best route to its proper neighbors, after prepending its own AS number to the route. When the best route at AS v is withdrawn due to some network failure event by the neighbor from where the route is learned, AS v will choose an alternative best route among the candidate routes and propagate the new best route to the proper neighbors. Note that an AS can only withdraw a previously announced BGP route. Irrelevant BGP withdrawal updates will be ignored by the receiving BGP routers. Therefore, from the viewpoint of security, it is sufficient to monitor BGP route announcements.

2.2 Internet Numbers Allocation and Assignment

Internet numbers, including both ASNs and IP addresses, are managed and allocated in a hierarchical fashion, coordinated by the Internet Assigned Numbers Authority (IANA) [12]. Figure 1 illustrates the Internet numbers allocation and assignment structure. IANA allocates numbers to the Regional Internet Registries (RIRs). Currently there are five RIRs, namely, AfriNIC, APNIC, ARIN, LACNIC, and RIPE NCC, which are responsible for the allocation of numbers in different regions of the global world. RIRs allocate Internet numbers to Local Internet Registries (LIRs) and National Internet Registries (NIRs). Internet Service Providers (ISPs) obtain Internet numbers from LIRs or NIRs or directly from RIRs. Users normally obtain Internet num-

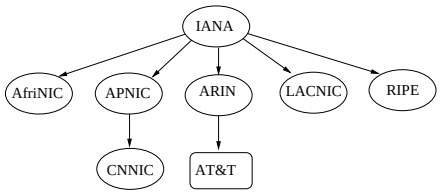

Figure 1: Internet numbers allocation and assignment structure.

bers from ISPs. In Figure 1, CNNIC and AT&T are example NIR and ISP, respectively. RIRs have well-established procedures for their members to request new Internet numbers.

Each RIR maintains a file of the Internet numbers that are allocated to its corresponding members [1]. We refer to this file as the numbers allocation file, or simply allocation file when there is no confusion. This file is updated and archived on a daily basis (midnight local time) and mirrored by all the RIRs. This file contains the most up-to-date Internet numbers allocation and assignment information for the corresponding RIR, with the following (partial) record format [2]:

```
registry|cc|type|start|value|date
```

where `registry` is one of the five RIRs, `cc` is the ISO 2-letter country code. `registry` and `cc` specify which RIR allocated the corresponding number and to which country the number was allocated or assigned, respectively. `type` can be asn (for AS numbers), ipv4 (for IPv4 addresses), or ipv6. It specifies the type of the record. In this paper we ignore IPv6 addresses. `start` and `value` specifies the first Internet number and the total count of elements in the allocated range, respectively. For example, if the record is for IP address allocation, `start` will contain the first IP address in the allocated range, and `value` indicates the total number of IP addresses in the range. Note that an allocated IP address range may not be represented by a single prefix. For example, an RIR may allocate a contiguous block of IP addresses to an NIR but the block cannot be specified by any single prefix. For this reason, we refer to an allocated IP address range as an IP address block when we discuss the IP address allocation from RIRs. `date` records the date when the allocation or assignment was made.

From the Internet numbers allocation records maintained by RIRs, we can obtain the country and RIR-level region to which an ASN and IP address block were allocated. We note here that, ARIN also maintains the information of certain legacy ASNs and IP address blocks that were allocated before RIRs were established. A number of these legacy ASNs and IP address blocks were allocated to countries that are not within the region of ARIN. We further note that although the allocation records are updated daily by all the RIRs, the file may contain outdated information. However, it is more accurate and complete than the voluntary registration information maintained by Internet Routing Registries (IRRs) [13], commonly known as `whois` databases, where the information is supplied by individual members and there is no formal and consistent method to verify the correctness of the information [27]. For this reason we only rely on the Internet numbers allocation information in our experimental studies; we do not use any IRR registration information. RIRs have been archiving the daily Internet numbers allocation file using the above format since late 2003 [1].

3. REGION-BASED BGP ANNOUNCEMENT FILTERING

In this section we first present a high-level overview of the region-based BGP announcement filtering (RBF) without considering practical deployment issues, for example, how an RBF-enabled BGP router obtains the ASN and IP address allocation information. The objective of this high-level overview is to illustrate how RBF can be effectively supported in the future Internet, which incorporates RBF-like schemes as a first-order design principle. We then consider the practical deployment of RBF in the current Internet.

3.1 RBF Overview

RBF can be deployed independently by individual ASes. An RBF-enabled AS will maintain the corresponding region-level Internet numbers allocation information. In this paper, we consider two levels of regions: the country-level and RIR-level. In the country-level RBF, an AS needs to maintain the country-level allocation information of all active ASNs and IP address prefixes, i.e., to which country an ASN (and IP address prefix) is allocated. Similarly, an AS needs to maintain the RIR-level allocation information of Internet numbers if it supports RIR-level RBF. An AS decides to support country or RIR-level RBF based on its own security requirements. A country-level RBF provides higher security guarantees than an RIR-level RBF. This allocation information can be maintained at a central server in the AS so that the border BGP routers can query. For better performance, all border BGP routers can also store a local cache of this allocation information.

When a BGP router receives a BGP announcement r, it will extract the origin AS v and the corresponding destination prefix p from the route r. The router will then look up the region information of v and p. If both belong to the same region, the route is accepted by the router for further normal BGP processing. If the origin AS and prefix do not belong to the same region, the router has identified a *mismatch* and the corresponding BGP announcement is *flagged* as a potential prefix hijacking attack. A flagged route is either dropped or passed to network operators for further analysis using techniques such as [10]. After RBF has been widely deployed on the Internet, we recommend the policy of dropping flagged routes for stronger security and better incentives for others to follow the operational practice consistent with RBF.

RBF is a light-weight prefix hijacking prevention and deterrence technique. From the above discussion, we can see that a BGP router only needs to maintain a small amount of Internet numbers allocation information (see Section 4 on the current counts of ASNs and

IP address blocks). The region lookup overhead should also be negligible using techniques such as hash tables (see the next subsection on the details of practical region lookup). In addition, RBF provides strong incentives for individual ASes to independently deploy the scheme. Even if only a single AS v deploys the scheme, it still guarantees that v's own traffic to a destination prefix will be sent to the correct destination (if the AS has a valid route to the destination).

So far, we have focused on the deployment of RBFs by individual ASes. RBFs can also be deployed at different scopes in a coordinated manner. For example, the *United States (US)* may require all its ISPs connecting to the external world (outside the US) to deploy RBFs. In particular, all such RBFs will filter out BGP announcements received from external world but involving network prefixes belonging to the US. In this way, US can at least guarantee that intra-US traffic will not be hijacked to a third-party country.

3.2 RBF Practical Deployment

In this section we discuss the practical issues we face in deploying RBF on the current Internet, in particular, how Internet numbers allocation information is obtained and maintained, the handling of sub-prefix announcements and prefix aggregation, and how to reduce the region lookup overhead of RBF.

3.2.1 Internet Numbers Allocation Information

As discussed in Section 2, each RIR updates and publishes its Internet numbers allocation information each day, and all the RIRs mirror each other's allocation information. An RBF-enabled AS can retrieve the numbers allocation information from the corresponding RIR it belongs to. However, an AS cannot obtain the real-time feed of the allocation information given that the allocation files are updated only daily. Therefore, it cannot access the allocation records made in the current day, and it has to rely on up to yesterday's allocation records to process today's BGP route announcements. In practice, this is not a problem. Given that it takes certain time for a network to obtain the ASN and IP address block, it is reasonable to assume that the RIR can add the corresponding allocation information into the numbers allocation file at least one day before the network is brought up online to the Internet. In this way, all the deployed RBFs on the Internet will have the correct allocation information of the new network (or rather the new ASN and IP address block).

3.2.2 Region Lookup of Prefixes

An RIR numbers allocation file frequently contains IP address blocks that cannot be represented by a single prefix. In addition, it is also common for an AS to announce a sub-prefix (smaller IP address range) instead of the original prefix allocated to the AS. Both contribute to the complexity of region lookup in RBF; simple hash tables may not be directly used. In the following we discuss a simple region lookup algorithm using a mechanism based on binary search. Other approaches are also possible.

We first note that the IP address blocks in the allocation files are non-overlapping (due to the fact that they are allocated to different organizations). Using this, we can sort all the allocation blocks by mapping the first allocated address in a block into a 32-bit integer. When a BGP announcement is received, the RBF-enabled router will extract the first IP address p_1 and the last IP address p_n from the announced prefix p, and convert them into 32-bit integers as well. The router then can find in the allocation set the largest integer n such that $n \leq p_1$, and verify that both p_1 and p_n belong to the corresponding block. Otherwise, the announced prefix will be flagged.

BGP routers may aggregate prefixes learned from downstream ASes into a single less specific prefix. In some cases, an *as_set* may be added into the *as_path* attribute due to prefix aggregation. *as_set* is an unordered list of ASNs, from which it is unclear which is the originating AS of a prefix. When an *as_path* contains an *as_set*, we will remove the *as_set* from the path attribute in determining the origin AS of the (aggregated) prefix. That is, we consider the first AS following *as_set* as the origin AS of the corresponding prefix. Given that prefix aggregation normally occurs for ASes with certain relationship, the involved ASes are likely within the same region. As a consequence, we believe that prefix aggregation will not have any major impacts on the performance of RBF when we follow the above approach to identify the origin AS of an (aggregated) prefix. We verify this observation in the next section when we study the performance of RBF.

3.2.3 Caching Legitimate ASN/Prefix Pairs

A BGP router may receive a large number of BGP announcements each day. It can be a prohibitive overhead for a router to look up the regions of the prefix and the corresponding originating AS for every BGP announcement it receives. In order to reduce the overhead of region lookup of RBF, a BGP router can maintain a cache table of the legitimate pairs of ASN and prefix. That is, after a router verifies that a pair of ASN and prefix observed in a BGP announcement belong to the same region, the router can add this pair into the cache table. After a BGP announcement is received by the router, it will first check if the pair of the prefix and the originating AS is in the cache table. The more expensive region lookup of the prefix and the ASN is only performed if the pair cannot be found in the cache table. Given that the majority of the association between prefixes and ASNs are stable (see the next section), we only need to look up the region for a small number of prefix/ASN pairs. This can greatly reduce the overhead of RBF.

4. PERFORMANCE EVALUATION

In this section we evaluate the performance of RBF using real Internet numbers allocation information and the BGP data trace collected by the University of Oregon Route Views project [21]. We will first examine the dynamics of the Internet numbers allocation information. We will then study the behavior of RBF under normal Internet operational conditions for which no

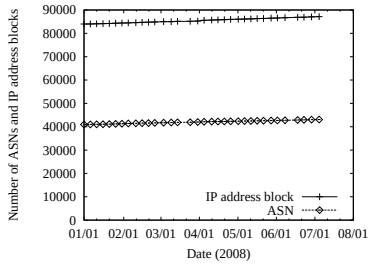

Figure 2: Number of total ASNs and IP address blocks.

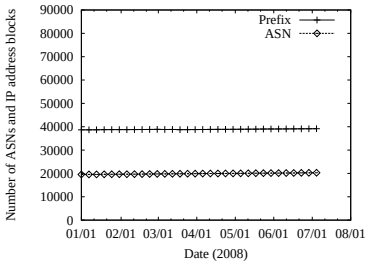

Figure 3: Number of total ASNs and IP address blocks (ARIN).

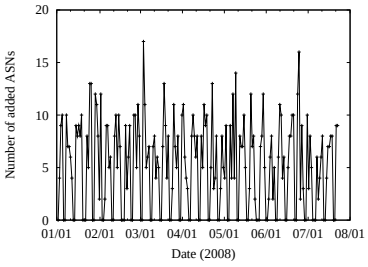

Figure 4: Number of added ASNs (ARIN).

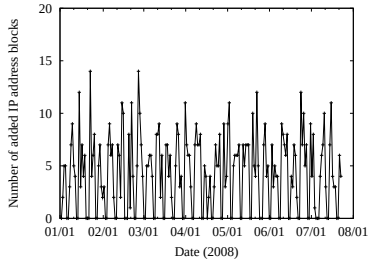

Figure 5: Number of added IP address blocks (ARIN).

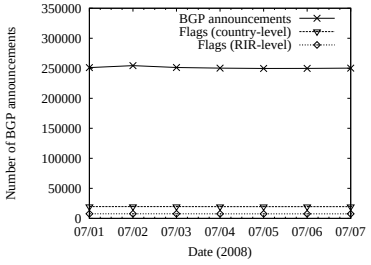

Figure 6: Flagged BGP announcements.

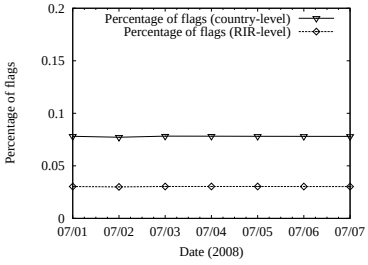

Figure 7: Percentage of flagged BGP announcements.

BGP prefix hijacking events were reported. At the end we will investigate the effectiveness of RBF on handling both large-scale and small-scale BGP prefix hijacking attacks. We summarize our main findings in the following.

- The Internet numbers allocation information is relatively stable. On average only a small number of (13) new ASNs and (19) IP address blocks are added daily on the Internet.
- About 8% (3%) of BGP announcements will be flagged by the country-level (RIR-level) RBF each day in normal Internet operations where no known hijacking events were publicly reported. The majority (98%) of flagged prefixes are announced from stub networks (i.e., edge of the Internet).
- Augmenting RBF with the information of legacy prefix announcement arrangement will dramatically reduce the number of flagged BGP announcements. On average only 215 BGP announcements will be flagged on a daily basis in the normal Internet operation. The flagged BGP announcements are relatively stable.
- Applying RBF to the well-documented AS9121 and YouTube incidents shows that RBF is very effective in detecting both large-scale and small-scale BGP prefix hijacking attacks.

4.1 Dynamics of Internet Numbers Allocation

In order to understand the dynamics of Internet numbers allocation information, we downloaded and examined allocation files from the five RIRs for six months and one week, from 01/01/2008 to 07/07/2008 (we use the week in July 2008 to study the behavior of RBF under normal conditions). Figure 2 shows the daily number of ASNs and IP address blocks during this period. On average, there are totally about $40\,K$ ASNs and $82\,K$ IP address blocks each day, combining all the five numbers allocation files. Importantly, the growth of the total ASNs and IP address blocks is stable and slow. On average, there are totally 13 new ASNs and 19 IP address blocks added each day. Among the five RIRs, ARIN and RIPE NCC allocate the most ASNs and IP address blocks (APNIC also contains a large number of IP address blocks, but relatively small number of ASNs). As an example, Figure 3 shows the daily number of ASNs and IP address blocks maintained by ARIN.

In order to provide a better view of the dynamics of the Internet numbers allocation, Figures 4 and 5 show the number of added ASNs and IP address blocks from ARIN, respectively. On average 6 ASNs and 4 IP address blocks were added daily over the studied period. We do not observe any deleted ASNs and IP address blocks in the period. Internet numbers allocation from other RIRs shows the similar trend (except that more ASNs were added than IP address blocks at other RIRs). From these figures we can see that the numbers of ASNs and IP address blocks are very stable, which indicates that relying on up-to yesterday's allocation information to filter the current BGP announcements will only in-

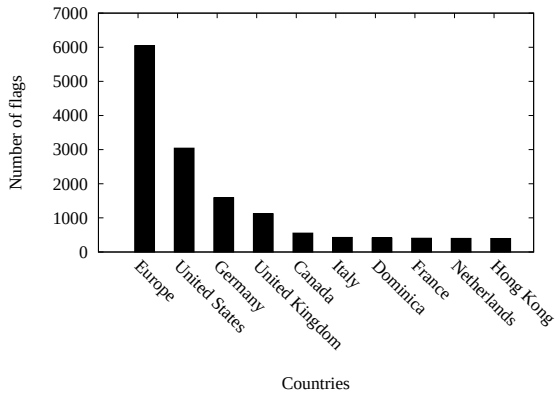

Figure 8: Top 10 countries announcing flagged prefixes.

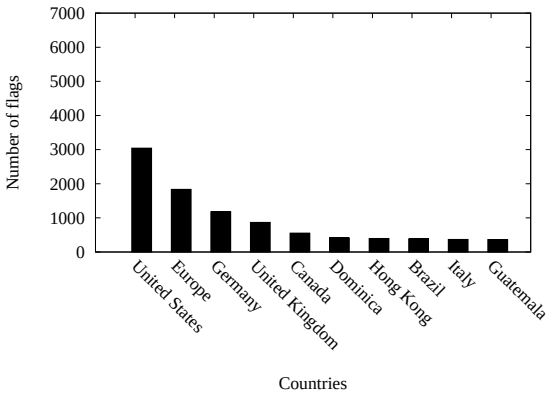

Figure 9: Top 10 countries announcing flagged prefixes (excluding EU).

cur a small error rate compared to using the real-time allocation information. In the following we will use the numbers allocation files in this manner when we study the RBF performance. Technically, we do have the allocation file for the current day since we are analyzing historical data. However, in order to study how RBF performs on today's Internet, we still use the yesterday's allocation file. Note that as we have discussed in the last section, when RBF is widely deployed on the Internet, RIRs should be able to publish the new allocation information at least one day ahead of the network being brought up online to avoid such errors.

4.2 RBF in Normal Internet Operation

In this section we study the behavior of RBF under normal Internet conditions where no prefix hijacking attacks were publicly reported. In order to detect BGP announcements with mismatched ASN and IP prefix allocation information, we download the BGP update messages from the Route Views project from 07/01/2008 to 07/07/2008. We use the previous day numbers allocation files from the five RIRs as the allocation information of ASNs and IP address prefixes for the current day. For example, in order to check if a BGP announcement is valid on 07/01/2008, we use the numbers allocation files updated at midnight 06/30/2008. As shown in the above subsection, this should only introduce a small error in the allocation mismatch detection. For each day, we pre-process the BGP announcements to remove duplicate BGP announcements.

4.2.1 RBF in Normal Internet Operation

Figures 6 and 7 show the daily number and percentage of flagged BGP announcements, respectively, for both country-level and RIR-level RBF. On average, there are about 250 K unique BGP announcements daily, of which, 19 K or 8% are flagged by the country-level RBF, i.e., the corresponding origin AS and IP prefix are in different countries. Among the announced BGP updates, 7 K or 3% are flagged by the RIR-level RBF.

From the above results we see that although the Internet numbers allocation information is updated daily by the RIRs, there are still a relatively high number of flagged BGP announcements. This is likely caused by two reasons. First, the allocation files may still contain a non-negligible amount of stale information despite daily update, due to, for example, unreported IP address prefix (and ASNs) ownership transfers. Second, flagged BGP announcements may also be caused by certain legacy peering relationship arrangements. For example, US ISPs announce a nontrivial number of prefixes allocated to other countries and regions. This is likely because originally US provides the Internet backbone service for a large portion of the Internet. We further note that, some of the flagged announcements may be related to prefix hijacks; not all hijacking events were reported publicly.

Towards the end of this subsection we will discuss how we can dramatically reduce the number of flagged BGP announcements under normal Internet operational conditions by augmenting RBF with a small hash table of the legacy prefix announcement arrangements. But first we perform an experiment to study the properties of mismatched BGP announcements. In particular, we are interested in learning which country and RIR regions announce more flagged prefixes and whose prefixes are more likely to be announced by an AS in a different region.

4.2.2 Distributions of BGP Flags

Figure 8 shows the top 10 countries that announced prefixes belonging to a different country, and the number of flagged BGP announcements from these countries. First we note that the country that announced most flagged BGP announcements is marked as "Europe". This is caused by the fact that, a large number of ASNs in the RIPE NCC allocation file are recorded as allocated to Europe instead of a specific country. From the figure we can also see that five European countries are ranked in the top 10 countries that announced prefixes belonging to a different country. We speculate this is likely caused by the fact that, European countries may provide Internet access for each other given that they are geographically close and tightly related in economy.

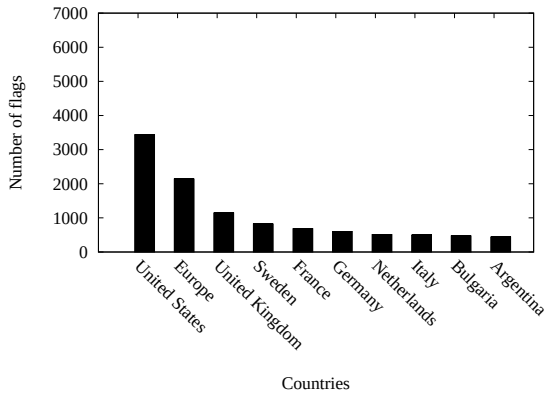

Figure 10: Top 10 countries whose prefixes announced by others.

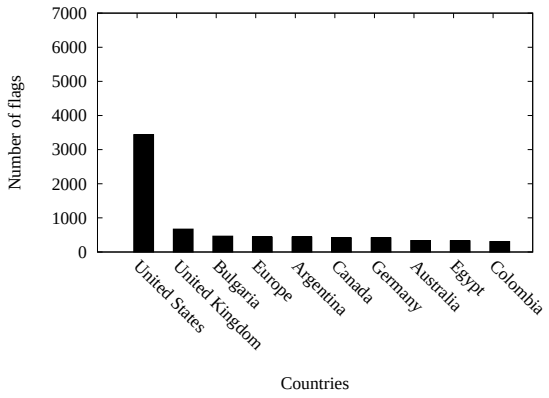

Figure 11: Top 10 countries whose prefixes announced by others (excluding EU).

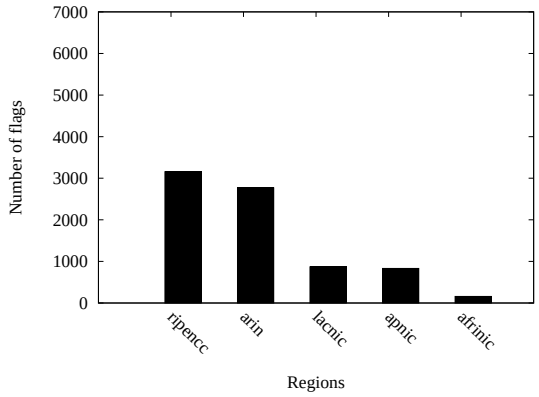

Figure 12: Regions ranked in number of flagged BGP announcements originated.

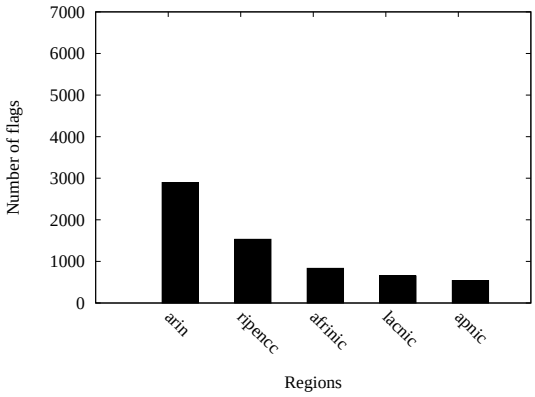

Figure 13: Regions ranked in number of prefixes announced by others.

To a degree, they should be considered as one country in BGP announcement filtering.

In Figure 9, we exclude the flagged BGP announcements whose origin AS belongs to "Europe" and the corresponding prefix belonging to an European country (or vice versa). After removing such cases, US becomes the top 1 country that announced prefixes belonging to a different country. In addition, two new countries become the top 10 countries in announcing prefixes belonging to a different country and two drop out of the top 10.

Figures 10 shows the top 10 countries whose IP address prefixes are announced by other countries. Figure 11 shows the corresponding results after excluding the flagged announcements whose origin AS is allocated to "Europe" and the announced prefix belongs to an European country (or vice versa). As we can see, the majority of the countries are European countries before excluding "Europe". This further confirms our above discussion that these countries may provide Internet access for each other due to close economic relationship.

Figures 12 and 13 rank the RIRs in terms of flagged BGP announcements originated by an RIR-level region, and the number of prefixes announced by a different region, respectively. From the figures we can see that RIPE NCC region announced most flagged BGP announcements, and ARIN is the region whose prefixes are most likely to be announced by other regions. These observations could be caused by a number of reasons. Historically, RIPE NCC provides Internet access service to some Asian countries. This may contribute to the result that RIPE NCC region originated most flagged BGP announcements. Second, ARIN maintains a large number of legacy prefix allocations. The information related to this legacy prefix allocation information may be outdated. These observations deserve further studies in order to fully understand the results and the history of the Internet development.

Figure 14 shows the total number of network prefixes an AS announces and the number of flagged prefixes announced by the AS, for country-level RBF. We only show the results for the ASes involved in at least one flagged BGP announcement. From the figure we cannot draw a clear conclusion on the trend of what ASes are more likely to announce flagged prefixes. ASes with all sizes (in terms of number of prefixes announced) may announce flagged prefixes, and smaller ASes appear to have a higher percentage of flagged prefixes than larger

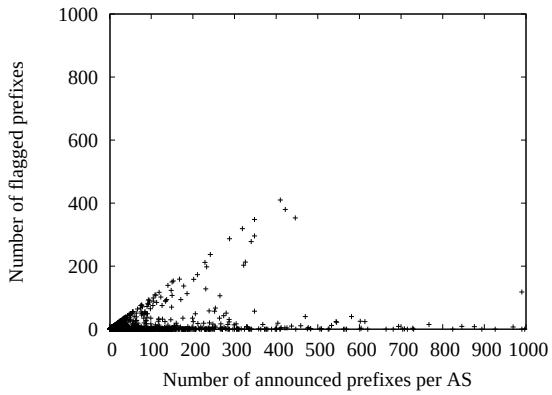

Figure 14: Distribution of flagged prefixes (country-level).

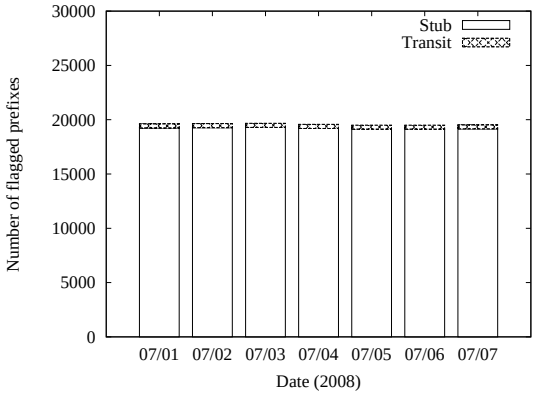

Figure 15: Network distribution of flagged prefixes (country-level).

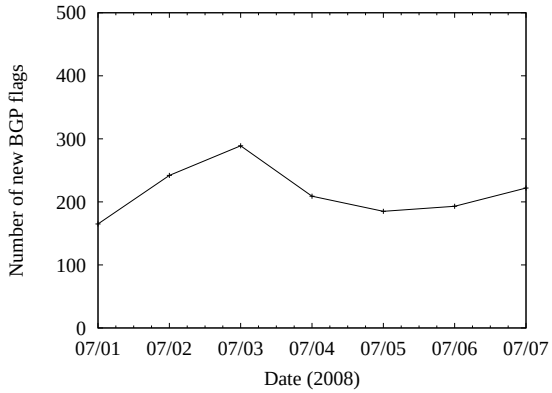

Figure 16: Number of flagged BGP announcements (country-level).

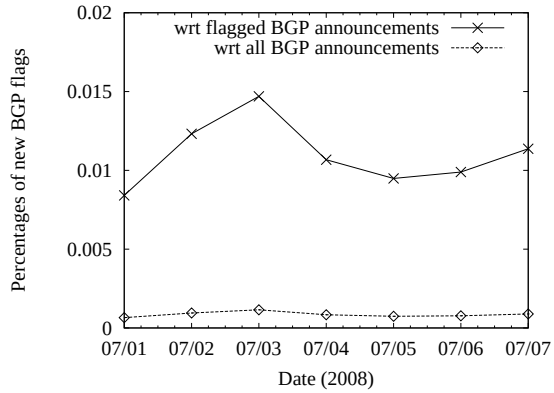

Figure 17: Percentage of flagged BGP announcements (country-level).

ASes.

In order to study which types of network domains are more likely to announce flagged prefixes, we classify networks into two types: stub and transit. Intuitively, stub networks are the edge of the Internet; they are customers of some ISPs and they do not provide Internet access service for any other ASes. In contrast, transit networks provide access service to others. We classify networks using the following heuristic. A network is considered as a stub network if its ASN only appears as the first (rightmost) ASN in the AS paths of the BGP announcements of prefixes [28]. Otherwise, it is considered as a transit network.

Figure 15 shows the daily number of flagged prefixes originated by different types of networks in the time interval (from 07/01/2008 to 07/07/2008), using the country-level RBF. As we can see from the figure, the vast majority (98%) of flagged prefixes were announced by stub networks. This is intuitively sound. First, stub networks are more likely to transfer IP address prefixes and ASNs to other stub networks following a company bankruptcy or merge. Second, from a security's viewpoint, transit networks are in general better managed and protected than stub networks. They are less likely to announce, i.e., hijack, prefixes belonging to a different AS (as a consequence of either misconfiguration or intentional attacks).

4.2.3 *Augmenting RBF with Legacy Arrangement Information*

In this subsection, we show that RBF can be augmented by a small mapping table of the legacy prefix arrangement that will trigger flagged BGP announcements. This augmentation can dramatically reduce the number of flagged BGP announcements. We note that we augment RBF in this way because we do not have the precise allocation information and a prefix may be announced by an AS in a different region due to historical reasons.

In order to obtain this mapping table for our experiment study, we obtain the BGP routing tables (and updates) on 06/30/2008, and consider all the flagged BGP announcements are caused by legacy ASN/prefix arrangement. That is, we treat all the flagged pairs of ASN and prefix on 06/30/2008 as valid announcements in the following days (from 07/01/2008 to 07/07/2008). We note that the decision to use the data on 06/30/2008

Table 1: Number of Flagged BGP announcements.

Date	Total BGP announcements	Flagged by country-level RBF (%)	Flagged by RIR-level RBF (%)
12/23/2004	143239	12985 (9%)	6455 (4%)
12/24/2004	248713	118172 (47%)	89534 (35%)

Table 2: Number of flagged prefixes announced by AS9121.

Date	Total prefixes announced	Flagged by country-level RBF (%)	Flagged by RIR-level RBF (%)
12/23/2004	55	0 (0%)	0 (0%)
12/24/2004	105530	105151 (99.6%)	83096 (78.7%)

is somewhat arbitrary (although no BGP prefix hijacking was reported on that day). In a real system, a sliding window with a few days' data should be used to establish the legacy prefix arrangement table.

Figures 16 and 17 show, for the country-level RBF, the daily number of flagged BGP announcements and the percentage of these flagged (with respect to the total number of BGP announcements and flagged BGP announcements without the mapping table augmentation), respectively. From the figures we see that on average only a small number (215) of BGP announcements will be flagged on a daily basis, which translates into 0.09% of the daily average total BGP announcements. Given this small number of flagged BGP announcements, network operators should be able to further analyze the causes and respond to the flags accordingly on a normal day.

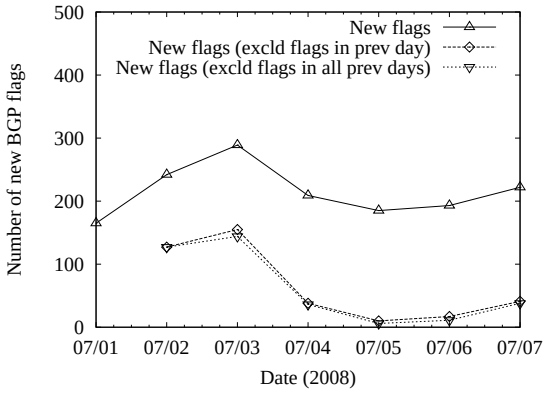

Figure 18: Number of new BGP flags (country-level).

In order to understand the dynamics of the BGP announcement flags, in Figure 18 we show the number of new BGP flags after excluding the flags seen in the previous day (*excld flags in prev day*) and the new BGP flags after excluding the flags seen in all the previous day since 07/01/2008 (*excld flags in all prev days*). Note first that removing old flags in these two ways results in almost identical number of new BGP flags. Second, on average, over 70% of BGP flags were seen in the previous days. That is, the flagged BGP announcements were relatively stable; the majority of them do not change frequently.

To further understand the nature of the ASes involved in the BGP flags, we obtained the degree of an AS, that is, the number of BGP peering neighbors of an AS. Figure 19 shows the degree of the ASes involved in the BGP flags on 07/03/2008, which has 289 BGP flags, the largest in the week (AS degrees on other days show the similar trend). As we can see from the figure, more than 60% of ASes have a degree of at least 5. In general, ASes with a large degree are normally service providers, which are less likely to announce false BGP routes. Flagged BGP announcements originated from service providers are more likely caused by prefix arrangement violating the policy of RBF, for example, providing Internet access service for networks in a different region (in this case, in different country).

We also classify the ASes into their corresponding countries, and show the top 10 countries that triggered BGP flags in Figure 20. In the figure we again show the country as "Europe" for the ASes whose country information shows "EU" instead of a specific country in the allocation file. In addition to the top 10 countries, we also cluster all the flags in the rest countries and show in the figure as "Other" (rightmost bar). We can see from the figure that the majority of the flags involve ASes from the European countries. In particular, 168 out of 289, or about 60% of BGP flags involve originating ASes and prefixes *both* from European countries. As we have discussed early, European countries provide Internet access for each other due to their geographical proximity and tight economic relationship. Such BGP announcements are likely to be valid ones but violate the RBF policy.

In summary, the flagged BGP announcements may be caused by new prefix announcement arrangement that violate the policy of RBF, or prefix hijacking. The way we use yesterday's allocation files to check the today's BGP announcements may also contribute to these mismatches. However, based on studies on the numbers allocation in Section 4.1, using yesterday's allocation files should only incur a small number of mismatches, if any. A more comprehensive study on the BGP flags is needed in order to understand the prefix announcement behavior of ASes and to improve the performance of RBF.

4.2.4 Impacts of Prefix Aggregation

In order to understand the impact of prefix aggregation on the performance of RBF, we also identified the ASes that aggregated prefixes. In this study we only consider prefix aggregation that results in *as_set* being carried in a BGP announcement. We consider an AS following the *as_set* as the AS performing prefix aggre-

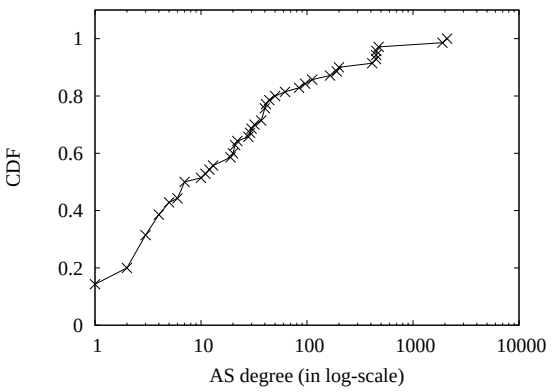

Figure 19: Degree of ASes involved in BGP flags (07/03/2008).

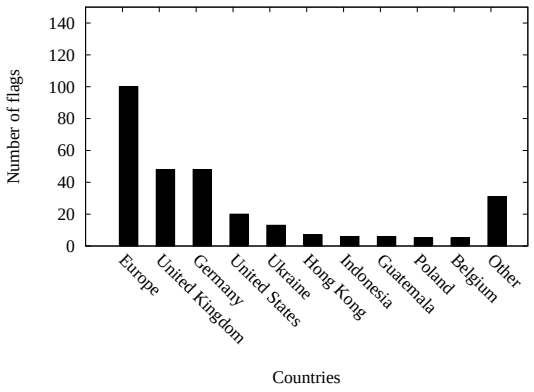

Figure 20: Country distribution of ASes involved in BGP flags (07/03/2008).

gation. Over the studied period from 07/01/2008 to 07/07/2008, the number of ASes that aggregated network prefixes is small ($21, 22, 22, 20, 20, 22, 22$ ASes on each day, respectively). Among these ASes, only two ASes triggered BGP announcement flags each day by RBF (for both country-level and RIR-level). Moreover, they are the same two ASes over the time duration. From the above observations, we conclude that prefix aggregation is only exercised by a small number of ASes, and importantly, it does not have any major impacts on the performance of RBF.

4.3 RBF with Known Prefix Hijacks

In this section we examine the effectiveness of RBF in detecting prefix hijacking attacks with known prefix hijack events. We examine the performance of RBF for both large-scale hijacking events involving a large number of network prefixes, and small-scale events that may only involve as little as a single network prefix. We use the following two well-documented prefix hijacking events as examples. The AS9121 hijacking event represents a large prefix hijacking event [22]. On 12/24/2005, AS9121 originated routes to about 100 K prefixes (about a full BGP routing table at the time). A large number of networks selected the routes originated by AS9121, attracting a large portion of Internet traffic destined to these prefixes to AS9121.

As an example of small-scale hijacking attacks, in the YouTube prefix hijacking attack [26], Pakistan Telecom (AS17557) originated prefix `208.65.153.0/24` belonging to YouTube. Because this prefix is more specific than the network prefix announced by YouTube (`208.65.152.0/22`), traffic to the YouTube servers on the `/24` address block was hijacked to AS17557, due to the longest prefix match algorithm used on the Internet for packet forwarding. Note that, as believed the original intention of Pakistan Telecom was to block the YouTube traffic *within Pakistan*; however, due to configuration errors this announcement was released onto the Internet, and affected the global access of YouTube. Because of the high popularity of YouTube (especially among young kids), the YouTube incident is considered as a notoriously bad example of prefix hijacking attacks, although it only involved a single network prefix. This incident showed to the broader community the weak security of BGP and what can be exploited to impact the normal operation of other networks.

We first examine the large-scale AS9121 prefix hijacking attack. Table 1 shows the number of total BGP announcements and flagged announcements by the country-level RBF and RIR-level RBF on 12/23/2004 (before the AS9121 incident) and 12/24/2004 (during the incident). From the table we see that on 12/23/2004, before the incident, the observations are consistent with what we observed in the previous subsection under the normal Internet operational conditions. There are about 9% of BGP announcements flagged by the country-level RBF and 4% by the RIR-level RBF. On 12/24/2005, the percentages of flagged announcements were increased to 47% and 35%, respectively, for country-level and RIR-level RBFs. Importantly, Table 2 shows the number of prefixes announced by AS9121, and that flagged by the country-level and RIR-level RBFs, respectively. From the table we see that AS9121 announced 55 prefixes on 12/23/2004, of which, none was flagged by the country-level and RIR-level RBFs. On 12/24/2004, AS9121 announced 105530 network prefixes, of which 105151 and 83096 prefixes were flagged by country-level and RIR-level RBFs, respectively. If country-level (RIR-level) RBFs have been deployed in an AS, it will filter out 99.6% (78.7%) of the hijacked routes.

For the YouTube hijacking incident, AS17557 started to announce the network prefix `208.65.153.0/24`, belonging to YouTube, on 02/24/2008. A country-level (RIR-level) RBF will successfully detect that the prefix is allocated to US (ARIN region) but the ASN 17557 is allocated to Pakistan (APNIC region). As a consequence, the BGP announcement will be successfully flagged and the hijacking attack is detected.

From the above examples (and intuitively from the operations of RBF), we can see that RBF is very effective in detecting both large-scale and small-scale cross-region prefix hijacking attacks, and confining all potential hijacking attacks to the same region. This helps

to provide better communications among ASes involved in prefix hijacking due to misconfiguration, and provide better deterrence to potential intentional prefix attacks because of the explicit legal consequence involved for hijacking prefixes from an AS in the same region.

5. DISCUSSION AND RELATED WORK

In this section we briefly discuss how RBF can be extended to detect path spoofing prefix hijacking attacks (in addition to the origin spoofing attacks). We also briefly summarize the related work.

5.1 Path Spoofing Prefix Hijacking

So far, we have focused on developing RBF to handle origin spoofing BGP prefix hijacking attacks. Here we briefly discuss how the allocation information may help dealing with path spoofing prefix hijacking attacks. As an example, we can classify network domains into stub ASes and transit ASes, where transit ASes provide transit services to others while stub ASes do not. We note that, stub ASes (and small regional ISPs) are less likely to suddenly provide transit service to another AS in a different (country or RIR level) regions. RBFs can flag such new BGP announcements as potential path spoofing prefix hijacking attacks. This is in line with the scheme proposed in [17], which requires two neighboring ASes to be geographically close. We will conduct a detailed study on detecting path spoofing prefix hijacking attacks in our future work.

5.2 Related Work

The work on improving BGP security can be broadly classified into two categories, those using cryptography and those do not. We discuss the ones using cryptography first. S-BGP [16] is the most comprehensive BGP security enhancement, which can prevent both types of BGP prefix hijacking attacks. However, S-BGP has not been widely deployed on the Internet due to the high management and packet processing cost, and the difficulty in identifying a widely accepted trust authority on the global Internet. soBGP [30], psBGP [29], and SPV [11] all tried to reduce the complexity of S-BGP (and eliminate the requirement of PKI).

Now we discuss the non-cryptography related work. The most related work is [27], which proposed to detect prefix hijacking attacks based on the ownership information of ASNs and IP address prefixes. The basic idea is that the destination prefix and the origin AS in a BGP announcement must belong to the same organization. However, the proposed scheme relies on the `whois` databases to determine if an ASN and prefix belong to the same organization. The `whois` databases contains both the allocation and voluntary registration information, and the registration information is much more incomplete and outdated than the allocation information. The authors proposed a number of heuristics in examining if an ASN and prefix belong to the same organization, for example, the `whois` records of the two have the same organization name, the same contact personnel, the same email address, the same DNS server, etc. In contrast, RBF only uses the numbers allocation information updated daily by the RIRs. Moreover, we propose the region-based BGP announcement filtering architecture as a way to facilitate the differentiated treatment of BGP announcements and out-of-band deterrence (legal consequence).

In pgBGP [14], a new BGP announcement will be quarantined for certain amount of time so that hijacking due to misconfiguration will not be propagated globally and operators are given longer time to diagnose the potential cause of the new announcement. Based on the observation that the association between a prefix and the originating AS and the peering relationship between neighboring ASes are reasonably stable, a learning-based approach was proposed in [23] to identify bogus BGP routes. These two schemes only rely on the control plane BGP routing information to detect prefix hijacking events. In contrast, RBF utilizes the AS number and prefix allocation information.

A light-weight distributed scheme was proposed in [32], which detects potential prefix hijacking attacks by monitoring data plane network distance changes to the target prefixes from a set of vantage points. The authors of [10] proposed to detect potential prefix hijacking attacks by data plane destination network fingerprinting. In iSPY [31] each network individually probes a set of transit ASes to determine its own reachability and infer if its own prefixes have been hijacked. These schemes use data plane probing to detect potential prefix hijacking events. RBF does not require any data plane probing.

PHAS [18] allows ASes to register their own prefixes and notices the registered ASes when possible hijacks occur. In IRV [8], each participating AS v will publish the prefixes announced by the AS (and the BGP announcement propagation information) in their DNS servers so that others can query if v has originated a prefix or announced a route to a neighboring AS. The scheme in [17] detects potential prefix hijacking attacks by exploiting the valley-free property of BGP routes and the geographical distance.

By correlating spam delivery with BGP announcements, [7] and [24] confirmed that spammers may hijack network prefixes to send spam in order to hide their identities. Hubble [15] is large-scale system to monitor the reachability problems on the Internet, in particular, the black holes where BGP routes exist to the corresponding prefixes but packets cannot reach the destinations, by periodically probing suspect network prefixes. Although it can be potentially used to detect prefix hijacking attacks, it focuses on a different problem than prefix hijacking. Campisano *et. al.* [4] developed a flow system based model to identify the root cause of individual BGP path changes. Similarly, this work does not directly address the prefix hijacking problem.

6. CONCLUSION

In this paper we developed a light-weight region-based BGP announcement filtering scheme (RBF) to improve the BGP security. In contrast to existing solutions that indifferently prevent or detect prefix hijacking attacks, RBF enables differentiated AS and prefix filtering treat-

ment and blends prefix hijacking prevention with deterrence. Experimental studies based on real Internet numbers allocation information and BGP traces showed that RBF is a feasible and effective scheme in improving BGP security.

Acknowledgment

We thank the anonymous reviewers of ACM ASIACCS 2010, whose invaluable comments and suggestions helped improve the quality and presentation of the paper.

7. REFERENCES

[1] ARIN. Allocated internet numbers. `ftp://ftp.arin.net/pub/stats/`.

[2] ARIN. RIR statistics exchange format. `ftp://ftp.arin.net/pub/stats/arin/README`.

[3] V. J. Bono. 7007 explanation and apology. `http://www.merit.edu/mail.archives/nanog/1997-04/msg00444.html`, Apr. 1997.

[4] A. Campisano, L. Cittadini, G. D. Battista, T. Refice, and C. Sasso. Tracking back the root cause of a path change in interdomain routing. In *IEEE/IFIP Network Operations and Management Symposium (NOMS 2008)*, Bahia, Brazil, Apr. 2008.

[5] H. Chan, D. Dash, A. Perrig, and H. Zhang. Modeling adoptability of secure BGP protocols. In *Proc. ACM SIGCOMM*, Sept. 2006.

[6] Cisco Systems, Inc. BGP path selection algorithm. `http://www.cisco.com/warp/public/459/25.shtml`.

[7] Z. Duan, K. Gopalan, and X. Yuan. Behavioral characteristics of spammers and their network reachability properties. In *IEEE International Conference on Communications (ICC)*, June 2007.

[8] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. McDaniel, and A. Rubin. Working around bgp: An incremental approach to improving security and accuracy of interdomain routing. In *Proceedings of Network and Distributed System Security Symposium*, San Diego, CA, February 2003.

[9] S. Halabi and D. McPherson. *Internet Routing Architectures*. Cisco Press, 2 edition, 2000.

[10] X. Hu and Z. M. Mao. Accurate real-time identification of IP prefix hijacking. In *Proc. IEEE Security and Privacy*, May 2007.

[11] Y.-C. Hu, A. Perrig, and M. Sirbu. SPV: Secure path vector routing for securing BGP. In *Proceedings of ACM SIGCOMM 2004*, Sept. 2004.

[12] IANA. Internet assigned numbers authority—number resources. `http://www.iana.org/numbers/`.

[13] IRR. Internet routing registry. `http://www.irr.net/`.

[14] J. Karlin, S. Forrest, and J. Rexford. Pretty good bgp: Improving BGP by cautiously adopting routes. In *Proceedings of IEEE International Conference on Network Protocols (ICNP)*, 2006.

[15] E. Katz-Bassett, H. V. Madhyastha, J. P. John, A. Krishnamurthy, D. Wetherall, and T. Anderson. Studying black holes in the internet with hubble. In *NSDI*, 2008.

[16] S. Kent, C. Lynn, and K. Seo. Secure border gateway protocol (S-BGP). *IEEE Journal on Selected Areas in Communications*, 18(4):582–592, 2000.

[17] C. Kruegel, D. Mutz, W. Robertson, and F. Valeur. Topology-based detection of anomalous BGP messages. In *6th Symposium on Recent Advances in Intrusion Detection (RAID)*, Sept. 2003.

[18] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang. PHAS: A prefix hijacking alert system. In *Proc. USENIX Security Symposium*, Aug. 2006.

[19] S. Murphy. BGP security vulnerabilities analysis. RFC 4272, Jan. 2006.

[20] O. Nordstrom and C. Dovrolis. Beware of BGP attacks. *ACM Computer Communications Review (CCR)*, 34(2), Apr. 2004.

[21] U. of Oregon. Route Views project. http://www.routeviews.org/.

[22] A. C. Popescu, B. J. Premore, and T. Underwood. The anatomy of a leak: AS9121. In *NANOG*, May 2005.

[23] J. Qiu, L. Gao, S. Ranjan, and A. Nucci. Detecting bogus bgp route information: Going beyond prefix hijacking. In *Proceedings of International Conference on Security and Privacy in Communication Networks*, Nice, France, Sept. 2007.

[24] A. Ramachandran and N. Feamster. Understanding the network-level behavior of spammers. In *Proc. ACM SIGCOMM*, Sept. 2006.

[25] Y. Rekhter and T. Li. A border gateway protocol 4 (BGP-4). RFC 1771, Mar. 1995.

[26] RIPE NCC. YouTube Hijacking: A RIPE NCC RIS Case Study. `http://www.ripe.net/news/study-youtube-hijacking.html`, Feb. 2008.

[27] G. Siganos and M. Faloutos. Neighborhood watch for internet routing: Can we improve the robustness of internet routing today? In *Proc. IEEE INFOCOM*, Anchorage, AK, May 2007.

[28] J. Stewart. *BGP4: Inter-Domain Routing In the Internet*. Addison-Wesley, 1999.

[29] T. Wan, E. Eranakis, and P. V. Oorschot. Pretty Secure BGP (psBGP). *ACM Transactions on Information and System Security*, July 2007.

[30] R. White. Securing BGP through secure origin BGP. *The Internet Protocol Journal*, Sept. 2003.

[31] Z. Zhang, Y. Zhang, Y. C. Hu, Z. M. Mao, and R. Bush. iSPY: Detecting ip prefix hijacking on my own. In *Proc. ACM SIGCOMM*, Seattle, WA, Aug. 2008.

[32] 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*, Kyoto, Japan, Aug. 2007.