IP Prefix Hijacking Detection Using Idle Scan

Seong-Cheol Hong[1], Hong-Taek Ju[2], and James W. Hong[1]

[1]Dept. of Computer Science and Engineering, POSTECH, Korea
[1]{pluto80, jwkhong}@postech.ac.kr

[2]Dept. of Computer Engineering, Keimyung University, Korea
juht@kmu.ac.kr

Abstract. The Internet is comprised of a lot of interconnected networks communicating reachability information using BGP. Due to the design based on trust between networks, IP prefix hijacking can occurs, which is caused by wrong routing information. This results in a serious security threat in the Internet routing system. In this paper, we present an effective and practical approach for detecting IP prefix hijacking without major change to the current routing infrastructure. To detect IP prefix hijacking event, we are monitoring routing update messages that show wrong announcement of IP prefix origin. When a suspicious BGP update that causes MOAS conflict is received, the detection system starts idle scan for IP ID probing so that distinguish IP prefix hijacking event from legitimate routing update.

Keywords: BGP, IP Prefix Hijacking, Routing, Security

1 Introduction

The Internet is a decentralized network comprised of many interconnected networks. Each network communicates reachability information using BGP (Border Gateway Protocol). The BGP is the de-facto inter-domain protocol that maintains a table of IP networks or prefixes. It designates network reachability among Autonomous System (AS) and there are more than 30,000 ASes in the Internet routing system [8]. The routers maintain and update their own routing table according to the routing information exchanged via BGP.

However, the Internet routing infrastructure is vulnerable to attacks due to lack of BGP security guarantees. The Internet was designed to provide communication on the basis of trust between networks. BGP also does not guarantee any security properties such as the authenticity of origin information and path attributes. IP prefix hijacking is the one of BGP security attacks, which a BGP router, for malicious purpose or by misconfiguration, announces an IP prefix that the router does not own. It results in reachability problem and communication failure in the Internet. IP prefix hijacking incidents are often reported on the NANOG mailing list [1].

To mitigate the impact of wrong routing information, some BGP extensions have been proposed, such as Secure BGP (S-BGP) [2], Secure Origin BGP (soBGP) [3]. Maybe these solutions can solve well-known BGP security problems, but it is difficult to deploy in practical network because the digital signature techniques that is used by S-BGP and soBGP cause high overhead and these improved protocols require changes to the existing protocol.

This work was partly supported by the IT R&D program of MKE/IITA [2008-F-016-02, CASFI] and WCU program through the KSEF of MEST, Korea [R31-2008-000-10100-0].

Many previous studies have proposed a method to detect IP prefix hijacking events [4, 5, 14, 15, 16, 17]. These are easily deployable solutions using passive monitoring or active probing. But some of these approaches use the routing registry information, such as IRR (Internet Routing Registry) databases, which can be outdated. The IP prefix hijacking events must be distinguished from legitimate routing updates because both cases cause MOAS (Multiple Origin AS) change.

To detect IP prefix hijacking event, we are monitoring routing update messages that show wrong announcement of IP prefix origin. When IP prefix hijacking occurs, there would be two networks having same IP address space in the Internet. Because the basic route selection process is to select routes with the shortest path, only the ASes close to the attacker AS are likely polluted. Our work focuses on fingerprinting two ASes having same IP prefix to distinguish IP prefix hijacking event from legitimate routing update. Our goal is to propose an easily deployable method that satisfies all of the following requirements: No modifications to routing protocol and current routers and performing the detection process without AS cooperation.

The organization of this paper is as follows. Section 2 describes the related works dealing with IP prefix hijacking. Section 3 defines the problem at hand and describes our solution approach. The experiment results are discussed in Section 4. Finally, Section 5 concludes the paper with possible future work.

2 Related Work

IP prefix hijacking is caused from an attack on the inter-domain routing protocol. The RPSEC (Routing Protocol Security Requirements) working group proposed a lot of Internet-Drafts about a scheme to improve routing protocol security, for examples, general security threats and requirements to routing protocols [6, 7]. Path attributes and Network Layer Reachability Information (NLRI) authentication is one of the requirements. This provides a means to verify and assure peering relationships and prefix advertisements against unauthorized announcements.

One of the BGP security architecture is S-BGP [2] which employs three security mechanisms – Public Key Infrastructure (PKI), optional BGP transitive path attribute and IPsec. S-BGP requires working with the Internet registries and ISPs to set up the PKT. However, PKI causes high overhead and requires a wide deployment in the Internet registries, router vendors and ISPs. The other proposal is soBGP [3] which is a deployable mechanism for validating the authorization of the BGP data. Its design goal is to be able to attain security profit without the participation of every AS and configure the level between security and overhead. While S-BGP and soBGP may be able to solve the security problems of routing protocol, it is not easy for both solutions to deploy in the current Internet infrastructure.

Zhao et al. [9] first explained MOAS conflict meaning that multiple origin ASes announce the same IP prefix. Originally, a unique AS number is allocated to each AS for use in BGP routing. However, many cases such as use of static routing and private AS number cause MOAS in the Internet [14]. When looking only the BGP update message, we cannot find any difference between legitimate MOAS and IP prefix hijacking.

Lad et al. [5] propose the method which monitors occurrence of new origin ASes in real time and notify the prefix owners that a suspicious update occurs. However, this method needs to rely on mutual cooperation between ASes. Karlin et al. [10] propose a system that automatically delays the use and propagations of suspicious routes. Introducing delay gives the human operators and systems time to investigate the suspicious route. The routers identify suspicious routes by consulting a table of trusted routing information learned from the recent history of BGP update messages. This method has some false positive cases which can legitimately occur: Provider change and occurrence of previously unseen provider.

In this paper, we propose a real-time detection method of IP prefix hijacking events. Our contribution in this paper is to propose an easily deployable detection method without AS cooperation.

3 Proposed IP Prefix Hijacking Detection Method

In this section, we propose a method to detect an illegal BGP update message. IP prefix hijacking events have some common characteristics such as MOAS and invalid route in a BGP message. Using these characteristics, we can identify problematic update messages and detect hijacking activities. In this study, we focus on the following objectives.

- Without changing BGP routing infrastructure
- Do not rely on mutual cooperation

The first objective means that the proposed detection approach must be easily deployable. The second requirement infers that the detection method should be effective without any AS cooperation.

3.1 IP Prefix Hijacking

IP prefix hijacking is a well-known security threat that corrupts the Internet routing tables. Each AS uses BGP to advertise its own prefixes to communicate with other ASes, but BGP does not provide any mechanisms to authenticate routing announcements. Therefore, a malicious router can announce wrong routing information to target prefix on the Internet without any authentication process. Sometimes, malicious users use IP prefix hijacking to get IP addresses on purpose to do spamming or DDoS attack.

IP prefix hijacking can occur on purpose or by accident in several ways. Many previous studies have classified IP prefix hijacking in detail [12, 16, 17]. We briefly explain the three types of IP prefix hijacking. Regular prefix hijacking occurs when the attacker AS announces a prefix that it does not actually own. As its wrong announcement is propagated, the Internet becomes to be polluted. Because the routers prefer the shortest AS path to forward traffic, not all of ASes in the Internet are polluted. Subprefix hijacking happens when the attacker AS announces a more specific prefix than what may be announced by the true origin AS. Most ASes are impacted by this announcement because the priority of more specific IP prefix is higher in route selection process. Lastly, IP prefix interception is that the attacker AS forwards the hijacked traffic to the origin AS. In this case, the victim cannot recognize the occurrence of prefix interception.

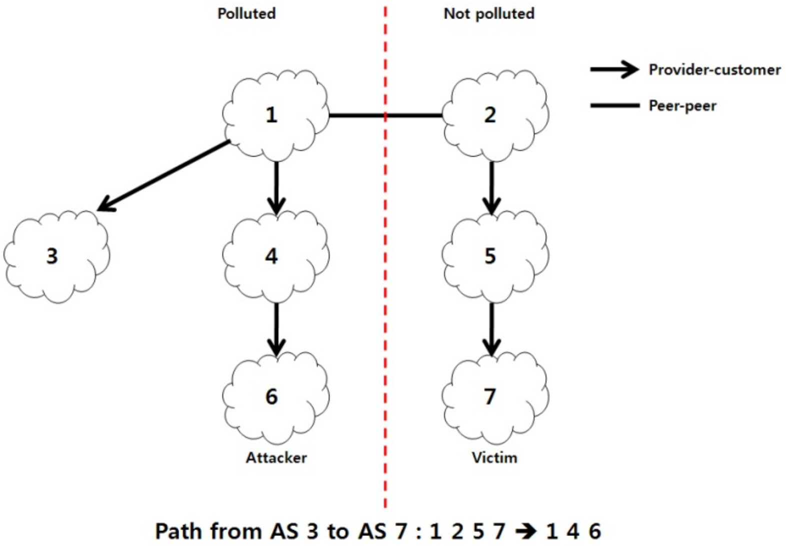

Fig. 1. Example of IP prefix hijacking: polluted and unpolluted ASes

Our approach focuses on the regular prefix hijacking. Fig. 1 shows an example of IP prefix hijacking with AS relationships. We suppose that the attacker AS is 6 and the victim AS is 7. When the attacker announces the IP prefix that the victim actually originated, this malicious routing information is propagated in the Internet. Typically ISPs can filter the announcements from their downstream ASes containing invalid IP address space, but previous hijacking incidents shows that it may not be applied by misconfiguration. With the given shortest AS path preference in routing, networks (AS 1 and AS 4 in Fig 1.) close to the attacker AS are polluted by the malicious announcement. AS 3 also receives the announcement and must decide whether the update is applied to routing table.

Because the routing tables of ASes near AS 3 are polluted, it cannot directly reach AS 7, but the unpolluted ASes can still arrive at the victim AS. A detection system using the information from multiple BGP monitoring points can recognize a MOAS conflict caused by IP prefix hijacking. However, it requires that monitoring points are located in both polluted and unpolluted ASes. That is, appropriate probing locations must be selected so that probing packets should reach two conflicted origin ASes through the different AS paths.

As mentioned above, we focus on the objective that we should avoid multiple vantage points. This requires additional techniques to properly detect IP prefix hijacking. Single vantage point cannot find any difference between legitimate MOAS and IP prefix hijacking.

We design the IP prefix hijacking detection algorithm using idle scan technique. Our algorithm identifies a suspicious BGP update message and verifies whether it is the IP prefix hijacking event.

3.2 Approach

In this section, we describe our solution approach. Fig. 2 shows an overview of our detection system in deployment. The detection system connects to the BGP router in observer AS to monitor BGP update messages and its routing table.

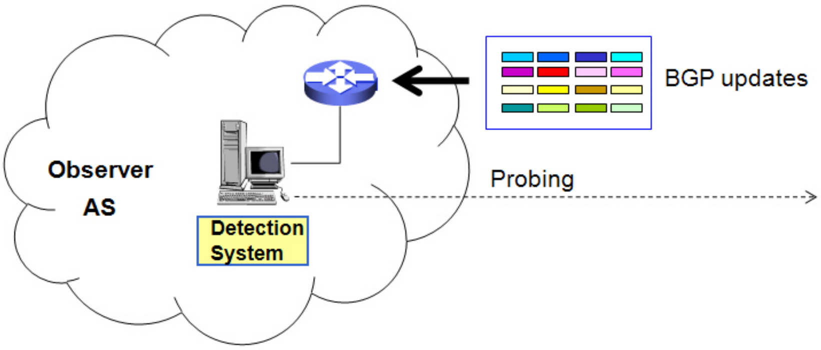

Fig. 2. Overview of the detection system in deployment

A BGP update message consists of withdrawn routes, the reachability information in the NLRI field and the AS_PATH attribute. The NLRI field indicates the IP address space about the destination AS and the AS_PATH attribute has the AS level path to reach the announced address space. With the comparison between update message and routing table, we can observe a suspicious event, especially MOAS conflict. When a suspicious BGP update that causes MOAS conflict is received, the system starts idle scan so that distinguish IP prefix hijacking event from legitimate routing update.

A probing technique called reflect-scan for fingerprinting the victim network is proposed [12]. This method is derived from the TCP idle scan technique described in [13]. The reflect-scan focuses on detecting subprefix hijacking cases, but it is applied to regular prefix hijacking in our approach.

The idle scan technique is used for completely blind port scanning that attackers can scan a target with sending a packet to a dummy host instead of the target. We utilize this technique on the purpose to reach the victim AS because we cannot directly arrive at the victim AS when IP prefix hijacking occurs as in Section 3.1. The key idea is to use the sequential IP ID increment property in IP packet and allow the unpolluted AS to forward the traffic to the victim AS.

The proposed technique is explained in Fig.3. First, the detection system selects a host (*Ha*) in the suspicious IP prefix, which satisfies the property to assign IP ID packets incrementally. Also, *Ha* should be idle because other traffic except for IP ID probing packets can interfere with the scan logic. And, the detection system should select a host (*Hr*) in the previous hop AS which is a just previous AS in the destination of AS_PATH to the target IP prefix. For example, if AS_PATH is ‘a b c d’, the previous hop AS is ‘c’ to the target AS ‘d’. This host should be alive and in service with open TCP port. The web server that always opens the HTTP port is a good candidate for *Hr*.

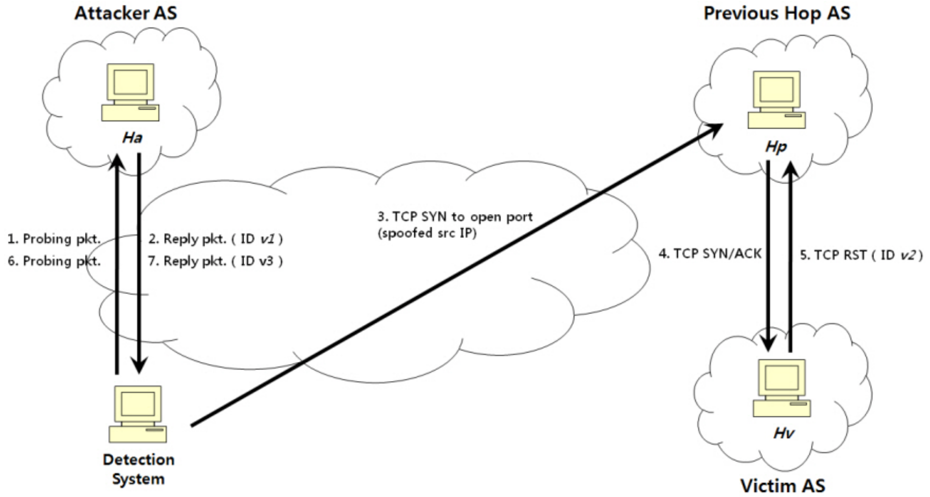

Fig. 3. Idle scan for IP prefix hijacking detection

After selecting the hosts, the detection system starts IP ID probing. The system sends a probing packet to *Ha* and records IP ID value in the reply packet. If a spoofed TCP SYN packet in which the source IP and *Ha*'s IP are same is sent to *Hr*, then *Hr* would response with a TCP SYN/ACK packet to *Ha*'s IP. When IP prefix hijacking occurs, *Ha* and *Hv* should be different. *Hv* that receives an unsolicited SYN/ACK packet will respond with a TCP RST. Therefore, one more probing to *Ha* can verify whether the received BGP update is the IP prefix hijacking event, because the IP ID difference between step 2 and step 7 is only one (that is, v3 = v1 + 1). In case of legitimate updates, *Ha* and *Hv* is the same host, and the IP ID difference is likely two or more (v3 = v2 + 1 = v1 + 2).

The probing packets used in step 1 and 6 do not need to be only TCP SYN/ACK packets like TCP idle scan technique. The proposed method requires the target hosts having predictable IP ID numbers for outgoing IP packets. To satisfy this requirement, we can select the protocol of probing packet that is expected to reply with incremental IP ID generation. More details are given in the next section.

The target hosts should be likely idle to reduce the false detection rate. To increase the detection accuracy, we can try to send multiple probing packets at step 1, 3, and 6. If the target is not as busy as well-known web server, we can sufficiently infer the occurrence of IP prefix hijacking as sending many probing packets.

4 Experiment Results

In this section, we present our experiment results to validate the proposed method and also discuss some of the obstacles.

We divide the validation process into three steps – correctness, feasibility and effectiveness. The correctness validation is that we test whether the proposed method detects the IP prefix hijacking events correctly. We should examine if the method can be used on real network, and the effectiveness of the proposed method should be measured from the performance point of view. In this paper, we carry out the correctness test and other validations are remained for future work.

4.1 Experiment Environment and Analysis

We performed an experiment to validate the correctness of our proposed method. Fig. 4 shows the experimental test-bed. We constructed the test-bed network which consists of routers and hosts using Linux machines. IP prefix hijacking condition is made by manipulating the routing table directly. The attacker and victim hosts can be operated on various operating.

This test-bed simulates the IP prefix hijacking case as described in Section 3.1. We suppose that Net6 is attacker AS and attempts to steal the IP prefix owned by Net7. After IP prefix hijacking event by Net6, Net1 and Net4 are polluted, but Net2 and Net5 are not polluted because the route selection process selects the shortest path. Fig. 4 shows that the routing tables are finally updated by the malicious action of Net6.

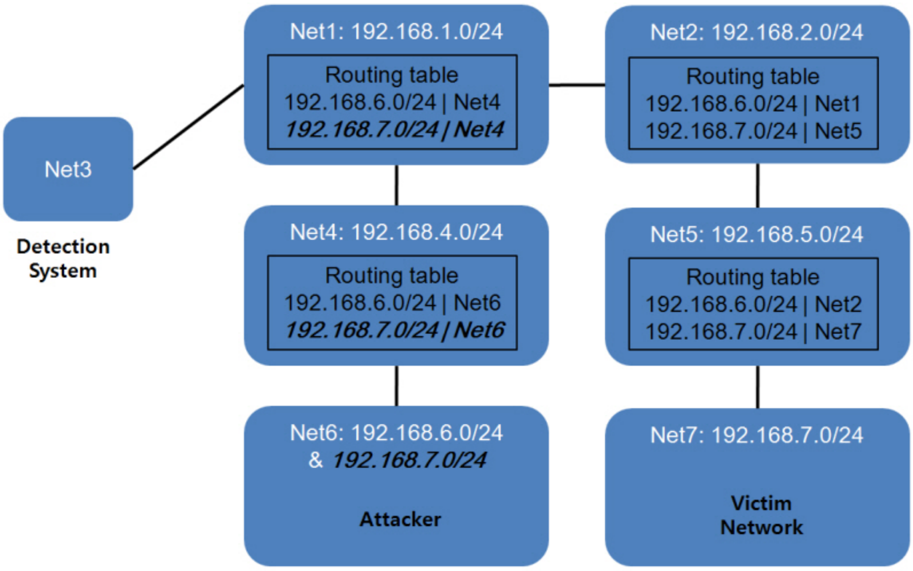

Fig. 4. Test-bed for correctness validation

The detection process is performed in Net3. We installed various operating systems (OS) on attacker AS Net6, such as Windows, Linux and etc., and checked for the IP ID probing process. As a result, the IP ID difference between the reply packets to two continuous probing packets was only one. Therefore, we conclude that idle scan technique is effective for detecting IP prefix hijacking events.

Table 1. IP ID generation pattern for different operating systems

Operating Systems	Reply Packets	
	TCP RST	ICMP
Windows	Incremental	Incremental
Linux	Zero	Incremental
Solaris	Incremental	Incremental
Router	Incremental	Random

Table 1 summarizes IP ID generation pattern of reply packets on various operating systems. We can use TCP probing for idle scan on the cases of most OS types except for Linux. Linux replies with zeroed IP ID packets in response to TCP SYN/ACK packets. In this case, we derive other protocols as reply packets, such as ICMP packets caused by UDP port scan to closed port. Linux also replies with sequential IP ID generation to ICMP packets. We can select appropriate probing packets to guarantee that the target systems assign IP ID packets incrementally on a global basis.

4.2 Discussions

When IP prefix hijacking occurs and an AS is impacted by that event, the AS cannot reach the victim network because the neighbor networks are already polluted. To distinguish a hijacking event from a legitimate update, we should be able to probe the attacker and victim AS. Idle scan can be a solution for that purpose and we utilize idle scan to detect IP prefix hijacking.

However, idle scan needs an appropriate target host having the property that the IP ID sequence generation happens incrementally. Therefore, we must perform port scan and OS identification on each AS in the routing table to find a candidate host. Another concern is that spoofing packets cannot be forwarded the target machine by egress filtering in ISPs. The higher service providers such as Tier-1 ISP are less likely to filter, so we can mitigate this problem by selecting an appropriate previous hop AS.

Also, the real network consists of a diversity of systems and devices, so we may suffer from unexpected responses in performing IP ID probing. We will solve this problem in future by feasibility tests in the Internet.

5. Concluding Remarks

In this paper, we have presented our algorithm of IP prefix hijacking detection. For detecting hijacking events, our algorithm relies on common characteristics of IP prefix hijacking such as MOAS and invalid route in BGP message. Our goal is to accurately identify hijacking events and distinguish them from the valid BGP updates. The proposed system does not require any protocol changes and multiple vantage points.

For future work, we will perform feasibility and effectiveness validation for the proposed method. The key problem is how to apply and validate the method in the real network. We will improve the algorithm to complement the reachability difficulties to the victim network.

We also plan to examine more IP prefix hijacking strategies and improve our algorithm that can detect various hijacking types. Furthermore, it is necessary to investigate the clear differences between IP prefix hijacking and valid updates.

References

1. North America Network Operator's Group, "NANOG: NANOG Mailing Lists," http://www.nanog.org/mailinglist/.
2. C. Lynn, J. Mikkelson, and K. Seo, "Secure BGP (S-BGP)," IETF Draft: draft-clynn-s-bgp-protocol-01.txt, June 2003.
3. Brian Weis, "Secure Origin BGP (soBGP) Certificates," IETF Draft: draft-weis-sobgp-certificates-02.txt., July, 2004.
4. C. Kruegel, D. Mutz, W. Robertson, and F. Valeur, "Topology-based Detection of Anomalous BGP Messages," In Proceedings of the 6th Symposium on Recent Advances in Intrusion Detection (RAID), LNCS 2820, Pittsburgh, PA, USA, September 2003, pp. 17-35.
5. M. Lad, D. Massey and D. Pei, "PHAS: A Prefix Hijacking Alert System," In Proceedings of the 15th USENIX Security Symposium, Vancouver, B.C., Canada, August 2006, pp. 153-166.
6. A. Barbir, S. Murphy, and Y. Yang, "Generic Threats to Routing Protocols," IETF Draft: draft-ietf-rpsec-routing-threats-07, October 2004.
7. B. Christian and T. Tauber, "BGP Security Requirements," IETF Draft: draft-ietf-rpsec-bgpsecrec-04, March 2006.
8. BGP Routing Table Analysis Reports, "BGP Reports," http://bgp.potaroo.net/.
9. X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. F. Wu, and L. Zhang, "An Analysis of BGP Multiple Origin AS (MOAS) Conflicts," In Proceedings of the 1st ACM SIGCOMM workshop on Internet Measurement, San Francisco, USA, November 2001, pp. 31-35.
10. J. Karlin, S. Forrest and J. Rexford, "Pretty Good BGP: Improving BGP by Cautiously Adopting Routes," In Proceedings of the 14th IEEE International Conference on Network Protocols, Santa Barbara, California, USA, November 2006, pp. 290-299.
11. H. Ballani, P. Francis, and X. Zhang, "A Study of Prefix Hijacking and Interception in the Internet," ACM SIGCOMM Computer Communication Review, Vol. 37, Issue 4, October 2007.
12. X. Hu and Z. M. Mao, "Accurate Real-time Identification of IP Prefix Hijacking," In Proceedings of the IEEE Security and Privacy, Oakland, California, USA, May 2007, pp. 3-17.
13. Insecure.Org, "TCP Idle Scan (-sI)," http://nmap.org/book/idlescan.html.
14. M. Tahara, N. Tateishi, T. Oimatsu, S. Majima, "A Method to Detect Prefix Hijacking by Using Ping Tests," In Proceedings of the 11th Asia-Pacific Symposium on Network Operations and Management (APNOMS 2008), LNCS 5297, Beijing, China, October 2008, pp. 390-398.
15. Z. M. Mao, J. Rexford, J. Wang, and R. H. Katz, "Towards an Accurate AS-level Traceroute Tool," In Proceedings of the 2003 Conference on Applications,

Technologies, Architectures, and Protocols for Computer Communications, Karlsruhe, Germany, August 2003, pp. 365-378.

16. C. Zheng, L. Ji, D. Pei, J. Wang, and P. Francis, “A Light-Weight Distributed Scheme for Detecting IP Prefix Hijacks in Real-time,” ACM SIGCOMM Computer Communication Review, Vol37, Issue4, October 2007.
17. Z. Zhang, Y. Zhang, Y. C. Hu, Z. M. Mao, and R. Bush, “iSPY: Detecting IP Prefix Hijacking on My Own,” In Proceedings of the ACM SIGCOMM 2008 conference on Data Communication, Seattle, USA, 2008, pp. 327-338.