Understanding IP Prefix Hijacking and its Detection

Christian Horn
(christian.horn@mail.tu-berlin.de)

Seminar Internet Routing,
Intelligent Networks (INET),
Research Group of Prof. Anja Feldmann, Ph.D.,
Technische Universität Berlin

June 8, 2009

Abstract

Since IP Prefix Hijacking is a major threat for every Autonomous System in the Internet, this paper tries to give an understanding of IP prefix hijacking and some of their detection methods. This may rise attention and awareness for that topic among the readers. If a malicious attacker would hijack an IP and use it for committing serious crimes, the original owner of the IP address would eventually be punished for that. Since with the given structure of routing IP Prefix Hijacks cannot be avoided, everybody should know what the potential problems are, so that they may be avoided in future designs.

1 Introduction

In February 2008 the Government of Pakistan and its Pakistan Telecommunication Authority (PTA) has directed the country's ISPs to block access to the website `www.youtube.com` to their users [15]. *On Sunday, 24 February 2008, Pakistan Telecom (AS17557) started an unauthorised announcement of the prefix 208.65.153.0/24. One of Pakistan Telecom's upstream providers, PCCW Global (AS3491) forwarded this announcement to the rest of the Internet, which resulted in the hijacking of YouTube traffic on a global scale* [14].

This incident shows that IP Prefix Hijacking is a real threat for every independent network known as Autonomous System (AS). IP Prefix Hijacking is an attack, where the IP prefix of a legitimate owner is hijacked, intentionally or by fault, through announcing a new route to the rest of the Internet. Today's Internet is an aggregation of many Autonomous Systems, which use the Border Gateway Protocol (BGP) to connect with each other. BGP is the *de-facto* standard inter-domain routing protocol to exchange routing information. Regarding the design of BGP [8] it is possible for an attacker to exploit the weaknesses of BGP, such as the lack of authentication. To carry out countermeasures the victims administrators have to detect the attack in the first place. Some work has already been done to develop techniques to detect IP prefix hijacking [2, 3, 6, 12]. Other work focuses on the improvement of BGP and the elimination of its weaknesses [16, 17], for instance the set-up of a global Public Key Infrastructure (PKI) with sBGP.

This paper is organized in the following way: In the first part IP Prefix Hijacking is described to gain fundamental knowledge of it. In the next section the techniques for detecting IP Prefix Hijacks are pointed out and reviewed. Afterwards the different techniques are discussed. A conclusion summarizes the paper.

2 IP Prefix Hijacking

The Border Gateway Protocol [7, 9] propagates routing information through the Internet and connects separate networks. It is a path-vector protocol, i.e. the complete path to a destination is exchanged in routing information updates between the routers.
With BGP it is not possible to verify the correctness of a prefix advertisement, since no authentication mechanisms are implemented in BGP to verify the origin. There is MD5 authentication between two peers, but no commonly used method to authenticate an AS Internet-wide. So it is easy for an attacker, once inside the "BGP network", to start his malicious work. To get access to the "BGP network" the attacker either has to gain full access to a BGP router, or another option would be having a connection to a provider that does not filter BGP traffic from its customers (like PCCW Global in the introductory example). Given that, the attacker could easily send false updates and rewrite routes to specific locations. This rewriting enables the attacker to harm the hijacked victim, commit crimes with the stolen identity or just intercept/modify the traffic. But not all IP Prefix Hijacks are done by malicious attackers. Much of the hijacks are done by mistake [10], like the unintentional misconfiguration of a BGP router by its administrators, where any victim can suffer from damages like in the case of a malicious attacker.

IP Prefix Hijacking can happen in various forms, but the key to it is always one (or more) compromised or poorly configured BGP router and the announcement of a false route for a victim's IP Prefix. Afterwards, these route spreads throughout the Internet. If the announced route is more appealing than an existing one, the decision would be to route traffic via the new route. Therefore every router that received this new, more appealing route should route its traffic destined to the hijack victim via the new route. Furthermore, to be stealthy and effective, the attacker should know the topology around its target and the target itself very well.

2.1 Attacking Techniques

The techniques for IP Prefix Hijacking are based on the announcement of invalid advertisements for IP prefixes, i.e. the BGP UPDATE message contains an AS-PATH, one of the important attributes in a BGP message, that is not the true AS-PATH to the victim's AS.

- **Hijack a complete prefix**
  The first way is to announce an AS-PATH that points to the attackers AS for the victim's IP prefix by defining it as the last AS in the path. This suggests that the attackers AS represents the victims prefix or has a direct link to the victim, causing other routers to send traffic to the attackers AS. This causes a Multiple Origin AS (MOAS) conflict [11], i.e. more than one AS claims ownership of an IP prefix. An example is that router *F* from *AS 4*, shown in figure 1, announces a route to a prefix possessed by *AS 1*, eventually causing *AS 5*, *AS 6* and *AS 7* to route the traffic destined for *AS 1* to *AS 4*. The second way is to announce a route which implies that the attackers AS is just in front of the origin AS. This avoids MOAS conflicts but enlarges the path by one AS. In that case *Router F* announces a direct route to *AS 1*, that does not exist in reality, so *AS 6* and *AS 7* may choose *AS 4* instead of *AS 5* for routing traffic to the IP prefix.
  For an attacker this technique is not the best choice, since it is not clear which one of the concurrent routes is chosen by BGP routers around the Internet. In the example *AS 3* and *AS 5* could still use the original route. So it would not be predictable how much

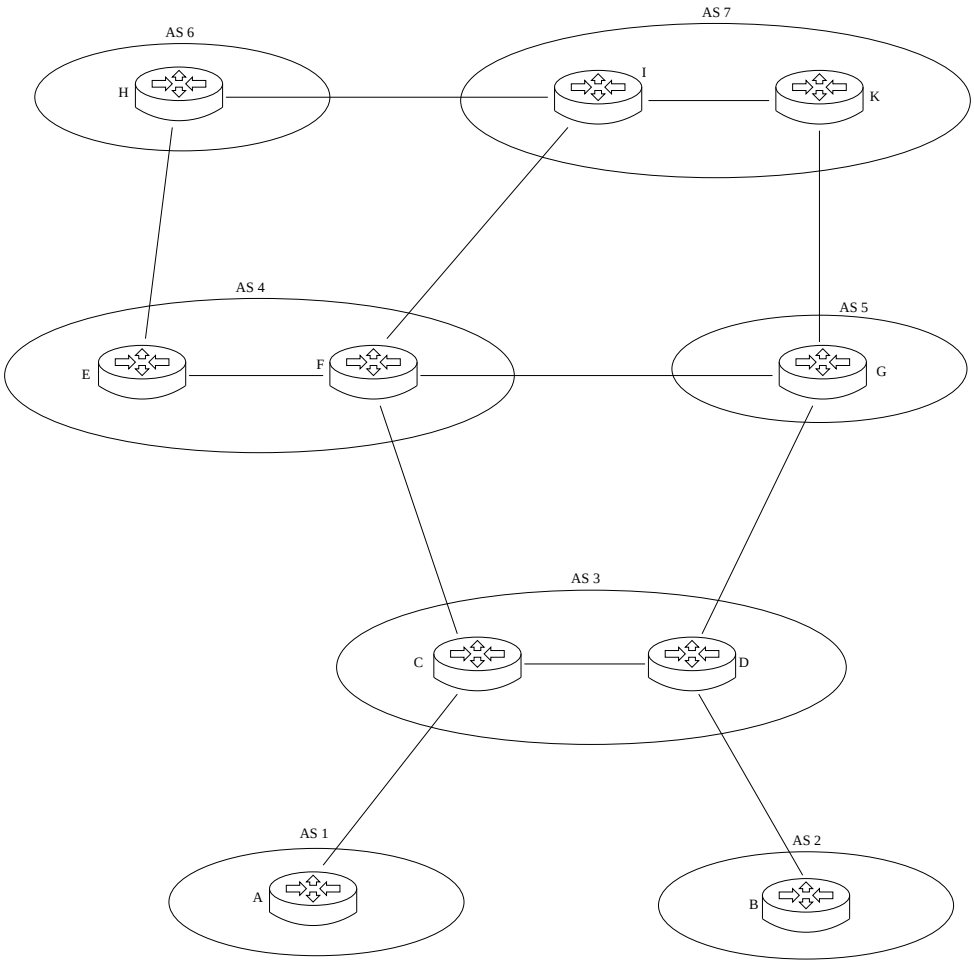

Figure 1: Example Autonomous Systems Topology

traffic is actually rerouted to the attacker, since routers near the original AS would still route the traffic to it, because of a shorter path. Even routers more far away could route the traffic to the original AS, because of its routing policy (e.g. cost efficiency or the like). An example would be *AS 7*, that uses *AS 5* because of a better/cheaper connection to it. Furthermore this can be easily detected, because of the two different routes to the IP prefix. To avoid the problem of concurrent routes, an attacker could send a WITHDRAWAL for the original route in a BGP UPDATE message. This would cause routers that receive the message to remove the original route from its tables. The problem here would be that the owner could detect that very fast and carry out countermeasures.

- **Hijack a sub-net of a prefix**
  As explained in the previous point, this can be done in similar ways. The difference is that the attacker announces a route to a subnet of the IP prefix that he wants to hijack. The advantage is that routers have the longest prefix match rule implemented, so the attackers route is always preferred by default. An example is that the original prefix would be 100.100.0.0/16 and the attacker announces 100.100.100.0/24. Depending on the use of the subnet by the original owner this can be exactly the traffic the attacker wants to intercept/modify. On the other hand if the subnet is not used by the owner, the attacker can use the IP address range to start malicious use of it without distracting

the traffic of the owner, thus maybe staying unnoticed by it.

- **Stealthy IP Prefix Hijacking**
  As shown in [5] there is also the possibility of using a stealthy attack. In that case

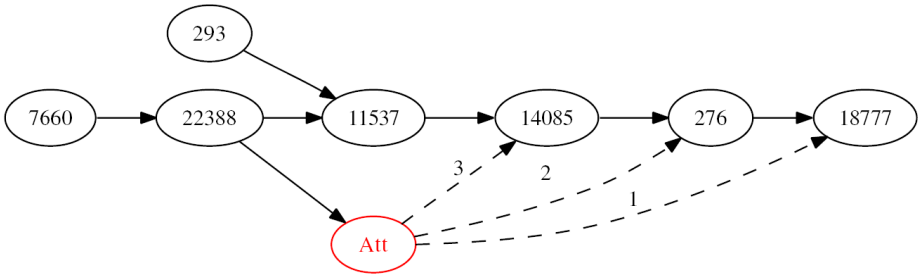

Figure 2: Stealthy Attacking Scenarios, taken from [5]

the attacker does not want to influence as many routers in as many ASes as possible, but only a small number. This can be done by announcing a slightly longer route than the original one to most routers, but not to all. So a small fraction of the traffic could be redirected nearly unnoticed. This, for instance, may ensure traffic recording over a longer period of time. *Figure 2 shows a partial representation of an AS tree for Texas State (AS 18777). It also shows three possible hijacking scenarios. The attacker establishes a peering relationship with AS 22388 and announces a path to Texas States prefix through AS 14085 (scenario 3), AS 276 (scenario 2), or AS 18777 (scenario 1). In scenarios 1 and 2, traffic from AS 22388 and AS 7660 destined to Texas State will be routed to the attacker. In scenario 3, since both the legitimate path and the attackers path have the same length, the effects are determined based upon local policies at AS 22388* [5]. Furthermore the authors of that paper show how to defeat current detection methods. For example they use a tool written by them called "fakeroute" to remain stealthy at all. Since some of the detection methods rely on traceroute, this tool can fake the responses by intercepting traceroute requests and falsifying its replies. The original route to the victim is recorded earlier and replayed later by fakeroute.

- **Hijacking by intercepting**
  There are two different scenarios for this technique. The first is to just intercept the traffic along its regular path. To be exact, this is no real IP prefix Hijacking, but results in the same consequences as described above. To do this, the attacker has to gain control over a router along the path and just intercepts/modifies or misuses the traffic coming by. In figure 1 the attacker just has to gain control over router *F* and log/modify all data coming by. No false BGP announcements have to be sent and no routing tables of other routers have to be influenced. Therefore, this kind of attack is much harder to detect. The second scenario is the hijacking of a prefix as described above, but leaving one of the neighbor ASes untouched, that is situated along the real path from the attacker to the victim. This means to announce a longer route to it by adding multiple instances of the attackers AS to the fake UPDATE message. Then the attacker can just forward the intercepted traffic to the original prefix owner via the untouched original route, making the prefix hijack harder to detect. In figure 1 for example router *E* was compromised and it wants to hijack a prefix of *AS 1*. *E* would announce an attractive fake route to *H* causing traffic from *ASes 6,7* to be redirected to *E*. To *F* the announced route would be enlarged by *E-E-E-E*, causing *F* to route its traffic for *AS 1* to *AS 3*.

2.2 Types of Attacks

As described in [3], [4] and [8], there are many different types of attacks. These attacks can be categorized into three groups, as done in [2]: Blackholing, Imposture and Interception. While all three types can have a different impact on the hijacked victim, the fact that the attacker has control of the traffic direction for the time-span of the attack stays the same. With a hijacked prefix the attacker is able to use the IP address range for other malicious activities like spamming [1], phishing and DoS attacks, without having the problem of covering the tracks after the attack. Any victim to such attacks would eventually blame the original owner of the IP address range. So it is important for any AS to detect the hijack in time to prove not being responsible for such attacks.

- **Blackholing**
  The attacker announces a false route to a machine that simply drops all packets intended for the victim or the announcement is an invalid IP address and all packets are going nowhere. This could be used for a DoS attack, where the regular Internet traffic for e.g. `www.google.com` could be redirected to the DoS victim that is eventually not prepared for such an amount of traffic and eventually crashes. Another reason for an attacker could be to simply cut off the IP prefix regular owner from the main part of the Internet and eventually causing financial losses for it. The reachability of the victim is affected either way.
- **Imposture**
  The attacker announces a false route to a machine that acts like the original server to e.g. steal information. That could be the cherry on the icing for a Phishing attack, because the URL would be the original one, so no client could distinguish the malicious server from the original one. In addition the original server is cut off from the traffic too, as in Blackholing.
- **Interception**
  The attacker announces a false route to a machine that forwards the traffic to the original server. This enables the classical man-in-the-middle scenario where the traffic could be intercepted, logged and modified and then forwarded to the original server. This kind of attack is much more difficult to detect than the other two, because the victim does not recognise any change in its amount of incoming traffic.

3 Detection of IP Prefix Hijacks

In [2], [3], [6] and [12] several approaches to detect IP Prefix Hijacks are described. Detecting a prefix hijack attack is not that easy, with a skilled attacker it could be nearly impossible. This is because of the vast size of the Internet, the lack of an administering authority and the closed policies of every AS. The present infrastructure and topology with BGP as the inter-domain protocol has its part too.

- **Monitoring IP Prefix Announcements** [12]
  The Prefix Hijack Alert System (PHAS) is one of the first approaches to detect IP prefix hijacking. The idea behind this approach sounds simple: monitor BGP routing data and report any announcement of a new AS associated with an IP prefix to the prefix owner. This could be done in several ways, but the paper uses the idea reporting via e-mail. It suggests to setup multiple email addresses with different servers in different IP prefix ranges, that, in the case of an hijack attack, one email would reach the prefix owner and not follows the hijacked path by mistake. Projects like the Route Views Project [13], which monitors the BGP traffic at several different locations around the Internet, could be used for that approach. This BGP update monitoring machines have to be set up to report the BGP changes to the alert systems servers, as shown in figure 3. Any change that is not authorized by the Prefix Owner would be detected in no time, if the email is checked automatically or manually in time.

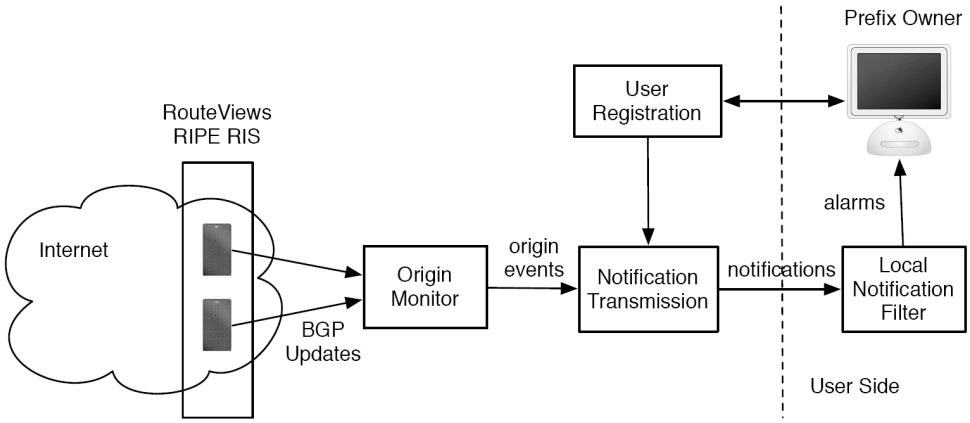

Figure 3: Components of PHAS, taken from [12]

- **Monitoring IP Prefix Announcements and Network Fingerprints** [3]

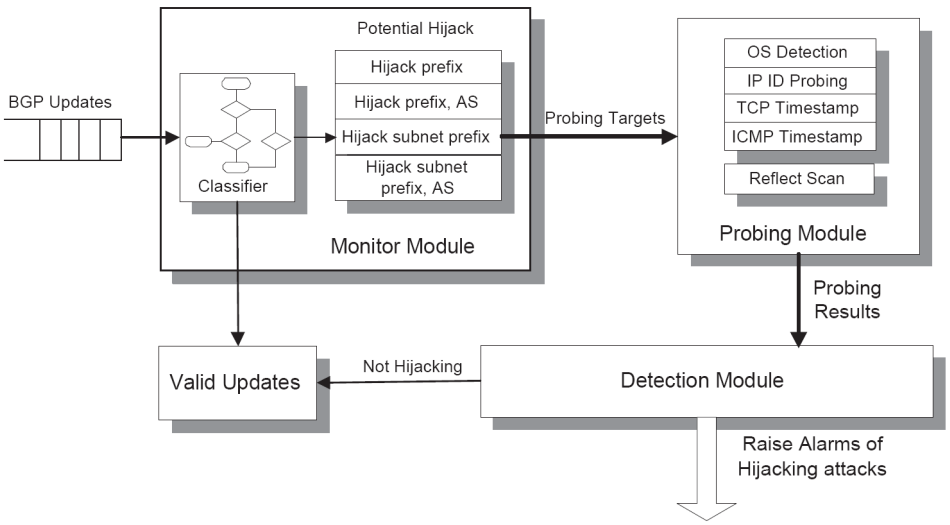

Figure 4: Architecure of detection system from [3],taken from [3]

The authors of that paper classified several attacking scenarios and made a detection system out of it. They use the monitoring of BGP routing data changes like in the previous approach extending it by probing the target network and hosts for special fingerprints. These fingerprints are generated using attributes like Operating System, host configurations, host software, host services, firewall policies, bandwidth information and characteristics of routers connecting the network, just to name a few. So every target network has a specific fingerprint that can be distinguished from other networks fingerprints. The fingerprints characteristics are collected using several techniques [3], for example port scanning patterns or IP ID probing, and tools like nmap. Figure 4 shows the architecture of that system. Once a BGP update arrives at the system, the classifier uses its data about known attacks to rank the update as valid or suspicious.

For the latter case, the destination network is probed and a fingerprint is generated. This fingerprint is compared to the stored one generated in the past. The detection module's analysis then distinguishes hijacking from valid updates.

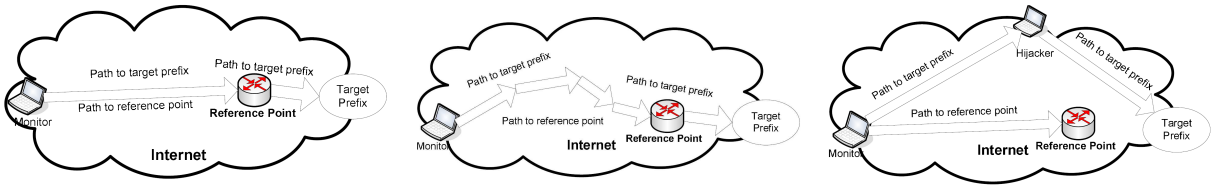

(a) Original (b) Legitimate route change (c) Prefix hijacking

Figure 5: Detecting Path Disagreement, taken from [2]

- **Monitoring from several Vantage Points** [2]
  This approach uses two observations: *1) the hop count of a path from a source to a destination is generally stable; and 2) the path from a source to a prefix is almost always a super-path of the path from the same source to a reference point along the previous path, as long as the reference point is topologically close to the prefix.* [2]. The vantage points or *monitors* try to detect a departure from these observations. This is done by monitoring the network location and path disagreement. The network location monitoring is done by measuring the hop-count distance from the vantage points to the destination network. If the hop count differs to a previous measurement and that difference is greater than a detection threshold, this is the first sign of a possible prefix hijack. The path disagreement measurement is done by observation of the path from one vantage point to its specific reference point. Every vantage point has its own reference point along the path from the vantage point to the target prefix. If the path to the reference point is no longer a sub-path of the path to the target prefix, this is the second sign of a possible prefix hijack. Figure 3 illustrates the path disagreement detection. There it becomes clear that the position of the vantage or reference point is a very critical aspect for that approach. Chosen at the wrong position the attacker could be situated between the reference point and the destination network. In that case nothing could be detected. To lower the possible error ratio, the measurements are done in parallel by several vantage points.
- **Probing of Transit ASes** [6]
  This is the single approach that can be used by an AS itself from inside its own network. The other approaches mentioned before need additional resources outside the AS to detect an IP prefix hijack. The authors of this paper propose a cyclic probing of the transit ASes, which interconnect the main parts of the Internet, from within the AS owning the prefix. The system aims to successfully probe at least one live IP per AS. This is done with a combination of traceroute, ping and TCP-ping. A database is maintained by the system that collects and updates live IPs and their related ASes. The idea behind it is that the traffic redirection through an IP prefix hijack attack would redirect the answers of the probing to the new bogus AS where the attacker is located. Thus the answer would never reach the origin, where the IP prefix belongs and the probes were sent. To avoid false alarms the received pattern of one cycle is compared to the results of the previous cycle and if a defined number of ASes is not reachable, or cut off, an hijacking alarm is triggered.

4 Discussion

These are the techniques that are available. Not one of them is able to detect IP prefix hijacks in an unfailing manner, since the structure of BGP, the vast size of the Internet and the lack of an administering authority complicates the detection. Every approach has its advantages and disadvantages, maybe a combination of all approaches could do the job.

Approach 1 suffers from inconsistent data for its analysis, since not the whole Internet's BGP routers can be monitored. In theory the BGP routers update each other, but nobody can predict the real impact on BGP routers Internet-wide in case of an attack. Using several monitors around the Internet can raise the probability of detecting an hijacking attempt, but the data inconsistency problem remains. The same applies to approach 2, and the extension of fingerprints is no magic bullet, since those characteristics can easily be faked by an attacker. Approach 3 sounds promising, but is highly dependent on the choice of the position for the vantage points. If the positions are chosen wrong, an attack can eventually not be detected. Since the "wrong" place depends on the position of the attacker, any place for a vantage point could be wrong. Approach 4 sounds promising too, but is not able to detect IP prefix hijacks that forward the traffic to the origin AS.

Another option to handle the problem of IP prefix hijacks is the improvement of BGP and the elimination of its weaknesses [16, 17]. With sBGP it could be possible to avoid hijacking problems, since sBGP demands the setup of a PKI for the whole Internet. So every BGP update could be authenticated in three ways: (a) through a certificate that ensures the allocation of an address space by the organisation advertised in the update, (b) through a certificate that authorizes a BGP router to speak on behalf of its AS and (c) through the use of IPSec and its authentication and integrity mechanisms. On the other hand the adaptation of BGP to sBGP is an expensive and complex step. This can be a reason for some organisations to avoid the introduction of sBGP in their systems.

5 Conclusion

This paper studied the IP Prefix Hijacking and its detection using several papers from different research groups. This made clear that IP prefix hijacking is a serious threat throughout the Internet and could affect everyone since none of the developed defending techniques have been deployed so far. Detecting IP prefix hijacks is difficult and not yet 100% reliable. A skilled attacker may remain undetected.

References

[1] A. Ramachandran, and N. Feamster. Understanding the Network-Level Behavior of Spammers. In *Proc. ACM SIGCOMM*, 2006.

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

[3] X. Hu, and Z.M. Mao, Accurate Real-Time Identification of IP Hijacking. In *Proc. IEEE Symposium on Security and Privacy*, 2007.

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

[5] C. McArthur,and M. S. Guirguis. Stealthy IP Prefix Hijacking: Dont Bite Off More Than You Can Chew. In *Proc. ACM SIGCOMM*, 2008.

[6] 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*, 2008.

[7] J. F. Kurose, and K. W. Ross. Computer Networking: A Top-Down Approach Featuring the Internet. Addison-Wesley, 2000.

[8] O. NordStröm, and C. Dovrolis. Beware of BGP attacks. *SIGCOMM Comput. Commun. Rev.*, vol. 34, no. 2, 2004.

[9] RFC 1771. `http://tools.ietf.org/html/rfc1771`

[10] P. Boothe, J. Hiebert, and R.Bush How Prevalent is Prefix Hijacking on the Internet. *NANOG36 Talk*, February 2006.

[11] 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 *Proc. ACM SIGCOMM*, 2001.

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

[13] University of Oregon Route Views Project. `http://www.routeviews.org/`

[14] YouTube Hijacking: A RIPE NCC RIS case study. `http://www.ripe.net/news/study-youtube-hijacking.html`.

[15] BBC News: Pakistan blocks YouTube website. `http://news.bbc.co.uk/2/hi/south_asia/7261727.stm`.

[16] L. Subramanian, V. Roth, I. Soica, S. Shenker, and R. Katz. Listen and Whisper: Security Mechanisms for BGP. In *Proc. First Symposium on Networked Systems Design and Implementation (NSDI)*, 2004.

[17] S. Kent, C. Lynn, J. Mikkelson, and K. Seo. Secure Border Gateway Protocol (S-BGP) – Real World Performance and Deployment Issues. In *Proc. IEEE NDSS Symposium*, February 2000.