

Evaluation of Prefix Hijacking Impact Based on Hinge-Transmit Property of BGP Routing System

Liu Yujing, Zhang Bofeng, Wang Fei, Su Jinshu
School of Computer, National University of Defense Technology, Changsha 410073, China
liuyujing@nudt.edu.cn
doi:10.4156/jnit.vol11.issue3.11

Abstract

BGP prefix hijacking is a sort of serious security threat of the Internet. In a hijacking attack, the attacker try to convince ASes to become infectors for redirecting data traffic to him in stead of the victim. The more infectors there are, the larger impact an attack has. It is important to understand the root of the matter why the impact of prefix hijacking differs a lot in different attacks. In this paper, by analyzing the BGP routing process under the control of routing polices and evaluating a series of Transmit factors, we realize that BGP routing system has a Hinge-Transmit property. It indicates that Tier-1 AS set is the hinge of the Internet, transmitting a large fraction of data traffic to the whole network; a subset of Tier-1 AS set with a special topological location (core AS set) is the hinge of data delivery paths to a specified destination, transmitting a large fraction of data traffic from any source to the destination. These hinge ASes are critical in transmitting large amount of data traffic in the Internet, and also critical in enlarging the impact of a prefix hijacking attack. From the aspect of Internet security, they should be protected from being infected carefully. Finally, we verify our findings by evaluating impacts of real hijacking incidents occurred in the Internet recorded by Route Views routing tables.

Keywords: BGP prefix hijacking, Hinge-Transmit property, Tier-1 AS, Core AS

1. Introduction

With the development of network technology, the Internet becomes a very important information source sharing by lots of users. However, as an infrastructure of the Internet, the inter-domain routing system is vulnerable to a variety of malicious attacks. BGP prefix hijacking is one sort of them. There were many prefix hijacking incidents occurred in the Internet, caused a large scale of damage in data reachability. In a hijacking attack, the attacker try to convince ASes to become infectors for redirecting data traffic to him in stead of the victim. The more infectors there are, the larger impact an attack has.

Previous efforts on BGP prefix hijacking can be sorted into three categories: preventions before the attack [1, 2], detections during the attack [3, 4] and reactions after the attack [5]. However, the impact evaluation of prefix hijacking is orthogonal to all the existing researches in the area. This topic is considered to be a valuable new start [6, 7].

Because the impact of BGP prefix hijacking has a tight relationship with BGP routing, we start with studying the property of BGP routing system. In this paper, we propose a method based on the fraction of amount of AS paths to evaluate the Transmit factor of every AS in the Internet. By evaluating Transmit factors of ASes wrt. the whole network and Transmit factors of ASes wrt. the specified destination, we realize that BGP routing system has a property of Hinge-Transmit. It indicates that Tier-1 AS set is the hinge of the Internet, transmitting a large fraction of data traffic to the whole network; a subset of Tier-1 AS set with a special topological location is the hinge of data delivery paths to a specified destination, transmitting a large fraction of data traffic from any source to the destination. If these critical ones receive the hijacking route and become infectors in a prefix hijacking incident, the impact of this attack will be enlarged significantly. It is much more crucial and effective to protect these hinge ASes from BGP prefix hijacking attacks. The evaluate results and conclusions in this paper have been verified by the BGP routing tables in Route Views project.

The rest of the paper is organized as follows. Section 2 reviews backgrounds on BGP routing and prefix hijacking. Section 3 presents a study of Hinge-Transmit property of BGP routing system.

Section 4 evaluates impact of prefix hijacking based on this property and shows evidence of our findings in real prefix hijack incidents. Section 5 makes a conclusion.

2. Background

In this section, we briefly review the process of inter-domain routing and prefix hijacking. It's much more complicated than other routing protocols because of the BGP routing policies that must be considered.

2.1. BGP Routing Policy

BGP enforces the routing policy of the Autonomous System, which corresponds to the business relationships with its neighboring ASes. There are three major types of business relationships between distinct ASes: provider to customer, customer to provider and peer to peer. The route selecting process depends on the export policies of upstream AS and the import policies of downstream AS.

According to the export policies, an AS usually does not transmit traffic between any of its providers or peers. When exchanging routing information with a provider or a peer, an AS does not export any routes of its providers and peers [8]. As shown in Fig. 1 (a), AS4 does not export routes of provider AS1 and peer AS2 to provider AS5 and peer AS6, but can export routes of any of its neighbors to customer AS7. According to the import policies, an AS applies a priority to every route learned from its neighbors. If a BGP router receives routes to a same destination from different neighbors, it prefers route from customer over those from peer and provider; then it prefers route from peer over that from provider [8]. As shown in Fig. 1 (b), AS5 receives routes to destination AS1 from provider AS2, peer AS3 and customer AS4. It chooses the route from customer AS4 as its best data delivery path.

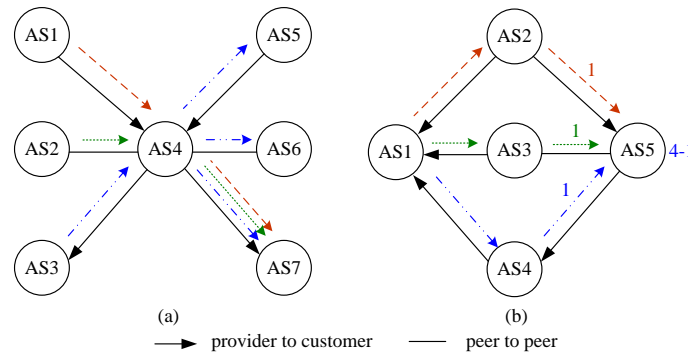


Figure 1. Examples of BGP routing policy

2.2. BGP Prefix Hijacking

Because of the lack of security mechanism, every BGP router has to believe the announcements they have received from other ones, no matter whether the message is creditable. As a result, the inter-domain routing system is vulnerable to misconfiguration and malicious attacks. Prefix hijack attack is a kind of serious security threat.

Before the attack, an AS announces its IP prefix to the Internet. The other ones who have learned this origin route can send data traffic to the origin AS in the future. As shown in Fig.2(a), AS1 announces 10.0.0/8 to its neighbors. With the propagation of this routing information, AS2, AS3 and AS4 get the route to the destination network (with IP prefix 10.0.0/8). During the prefix hijacking attack, the attacker announces IP prefix which belongs to the victim network. Such false hijacking route propagates on the Internet, too. The ones who choose to believe it become infectors. Data traffic from those polluted ASes will be sent to the attacker instead of the victim. Fig 2(b) illustrates this scenario. Attacker AS4 wrongly announces victim AS1's prefix. AS3 is infected as a result of

accepting the hijacking route as its best choice. Consequently, the data traffic from AS3 to AS1 will be redirected to AS4.

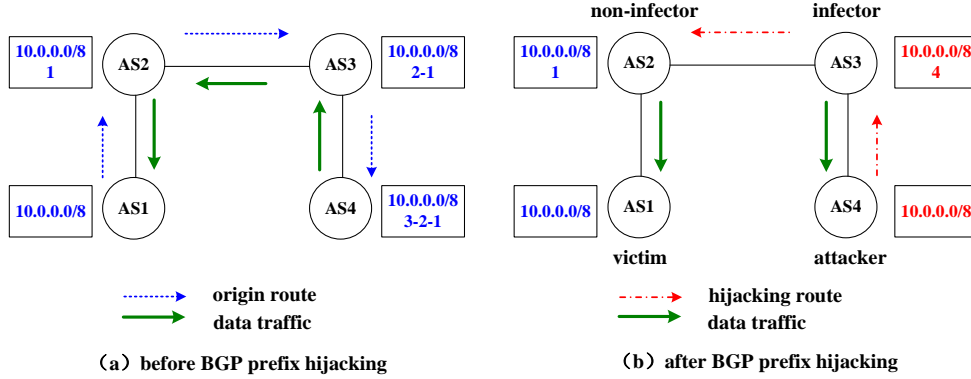


Figure 2. Examples of BGP prefix hijacking

3. Study of Hinge-Transmit Property

In this section, we define a measurement of AS denoting how crucial it is for transmitting data traffic, and find the Hinge-Transmit property of BGP routing system from our evaluation results. This property is verified by routing tables recorded by Route Views project.

3.1. Transmit Factor Definition

In BGP routing table, the route from source AS to destination AS is recorded as the AS path attribute in every entry as shown in Table 1. This example indicates that the data traffic from local network to the network with IP prefix 3.0.0.0/8 will be transmitted through AS100, AS200, AS300, and then reaches the destination AS400. AS100, AS200 and AS300 are the transit ASes which transmit traffic for local network in this instance.

Table 1. BGP routing table entry

Timestamp	Peer IP	Peer AS	Prefix	AS Path	...
1222825411	10.0.0.1	100	3.0.0.0/8	100 200 300 400	...

The more times a transit AS appears in the AS paths, the more traffic it is responsible for transmitting, the more critical it is in the inter-domain routing system. According to this evidence, we define the *Transmit factor* of an AS wrt. a destination AS: the fraction of AS paths from any source to the destination which contain this AS as a transit AS. It's used to describe how much traffic this AS transmits to the destination. Likewise, we define the Transmit factor of an AS wrt. whole network: the fraction of AS paths from any source to any destination which contain this AS as a transit AS. It's used to describe how much traffic this AS transmits to the whole network. Formally, in an AS set N consists of all ASes in the whole network, $AP_{s,d}$ represents the AS set consists of all the ASes contained in the AS path from source AS s to destination AS d . For every AS i in N , $E_{s,d}(i)$ equals to 1 if $i \in AP_{s,d} - \{d\}$. Otherwise, it equals to 0. The Transmit factor of i wrt. destination d is defined as (1); the Transmit factor of i wrt. whole network is defined as (2).

$$T_d(i) = \frac{1}{|N|} \sum_{s \in N} E_{s,d}(i) \quad (1)$$

$$T(i) = \frac{1}{|N|^2} \sum_{s \in N} \sum_{d \in N} E_{s,d}(i) \quad (2)$$

Furthermore, we define the Transmit factor of an AS set wrt. a destination AS: the fraction of AS paths from any source to the destination which contain at least one AS in this set as a transit AS. It's used to describe how much traffic this AS set transmits to the destination. Likewise, we define the Transmit factor of an AS set wrt. whole network: the fraction of AS paths from any source to any destination which contain at least one AS in this set as a transit AS. It's used to describe how much traffic this AS set transmits to the whole network. Formally, I represents an AS set. $E_{s,d}(I)$ equals to 1 if $\exists i \in I, i \in AP_{s,d} - \{d\}$. Otherwise, it equals to 0. The Transmit factor of AS set I wrt. destination d is defined as (3). The Transmit factor of AS set I wrt. whole network is defined as (4).

$$T_d(I) = \frac{1}{|N|} \sum_{s \in N} E_{s,d}(I) \quad (3)$$

$$T(I) = \frac{1}{|N|^2} \sum_{s \in N} \sum_{d \in N} E_{s,d}(I) \quad (4)$$

3.2. Transmit Factor Evaluation

In order to evaluate the Transmit factor of every AS accurately, it is essential to collect all the AS paths from any source to any destination. As mentioned in above paragraphs, BGP is a policy-based routing protocol. Based on the basic limits of routing polices, it is possible to infer all the AS paths by simulating the BGP routing process if AS topology with relationship information of the Internet is offered. In this section, the AS Relationships data in CAIDA [9-11] is used to infer the AS paths.

3.2.1 Transmit Factor Evaluation wrt. Whole Network

We use the inferred AS paths as an input and calculate the Transmit factor of every AS in the Internet defined in (2). The result is shown as Fig. 3. From this figure, we find that there are only a few ASes having very high Transmit factor, but most ASes having very low Transmit factor. This means in BGP routing system, a few ASes are responsible for transmitting a large fraction of data traffic to the whole network, but most ASes transmit few.




Figure 3. The result of Transmit factor evaluation wrt. whole network

Those ASes with high Transmit factor belong to the Tier-1 AS set [12]. These 13 Tier-1 ASes are in the top routing hierarchy of the Internet. They peer with each other and have no providers. We consider those ASes as one set called *Tier1S*, and evaluate the Transmit factor of this set defined in (4). The result of $T(\textit{Tier1S})$ is as high as 0.827, which means there are 82.7% of the total AS paths from any source to any destination passing through at least one Tier-1 AS. From this result, we realize that in BGP routing system, the Tier-1 AS set is the hinge of the Internet data delivery, transmitting a large fraction of data traffic to the whole network.

3.2.2. Transmit Factor Evaluation wrt. Destination

The evaluation result of Transmit factor of AS wrt. destination defined as (1) shows that even in the Tier-1 AS set, not all the ASes are equivalent. The Transmit factor differs sharply by different destinations. An example of our evaluation results is presented in Fig. 4. The Transmit factor of Tier-1 AS2914 wrt. destination AS32148 is high, which means a large fraction of traffic from any source to AS32148 is transmitted by AS2914; and the Transmit factor of Tier-1 AS3561 and AS209 wrt. destination AS25650 are high, which means a large fraction of traffic from any source to AS25650 is transmitted by AS3561 and AS209.

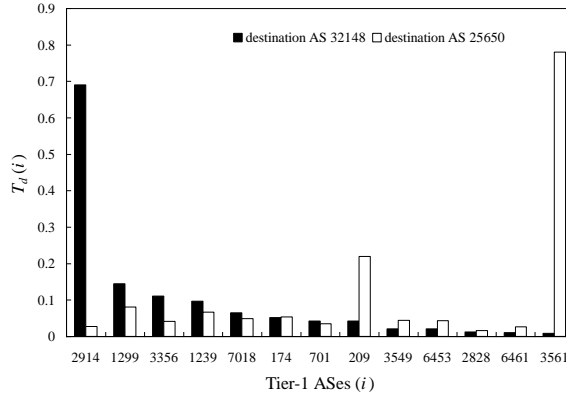


Figure 4. An example of Transmit factor of Tier-1 ASes differs by different destinations

By analyzing the topology of these ASes, we find these Tier-1 ASes with high Transmit factor are in special topological locations wrt. the destinations. As shown in Fig. 5, AS2914 is the Tier-1 AS which has the least provider - customer hops to the destination AS32148; while AS3561 and AS209 are the Tier-1 ASes which have the least provider - customer hops to the destination AS25650.

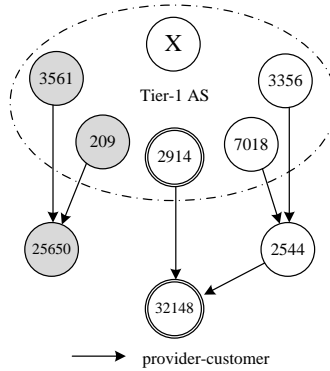


Figure 5. The topological locations of ASes with high Transmit factor wrt. destinations

We call the Tier-1 AS in this special topological location the *core AS* of the destination. In BGP routing system, almost every destination AS has a core AS set $core_d$ which consists of the nearest Tier-1 providers of the destination. To validate all the ASes' core AS sets have high Transmit factors wrt. destinations, we evaluate all of them defined in (3) and calculate the mean value defined as (5).

$$\bar{T}(I) = \frac{1}{|N|} \sum_{d \in N} T_d(I) \quad (5)$$

The result of $\bar{T}(core_d)$ is 0.652, which means there are 65.2% of the total AS paths from any source to a destination passing through at least one of its core AS. This result shows that every AS' core AS set transmits a large fraction of data traffic from any source to the AS, acting as the hinge of the data delivery paths to the destination.

3.3. Hinge-Transmit Property of BGP Routing System

This phenomenon is caused by the routing policy of BGP. As mentioned before, a BGP router prefers route from customer over others when having several alternate routes to reach a particular destination. So the routing information propagates along the least customer - provider hops from the destination to Tier-1 ASes, and then propagates along peer - peer or provider - customer hops after that. This causes many BGP routing paths to pass through the core ASes, which means there is a large fraction of data traffic transmitted by a small amount of ASes. We call this property of BGP routing system *Hinge-Transmit*. To be more convincible, we validate this property by BGP routing tables collected by Route Views [13, 14].

We choose routing tables in 10 different times randomly. By calculating the fraction of AS paths which contain Tier-1 ASes, we get the result shown in Fig. 6. The result shows that in all 10 cases, more than 80% of the AS paths pass through at least one Tier-1 AS. The Tier-1 AS set not only has a high connecting degree but also has a high delivery degree. Further more, we calculate the fraction of AS paths which contain core ASes with the help of the topology information offered by CAIDA AS Relationships project, and get the result shown in Fig. 7. It shows that more than 75% of the AS paths from any source to a destination pass through at least one core AS of the destination in all 10 cases. The core AS set plays a critical role in transmitting data traffic to the destination in BGP routing system.

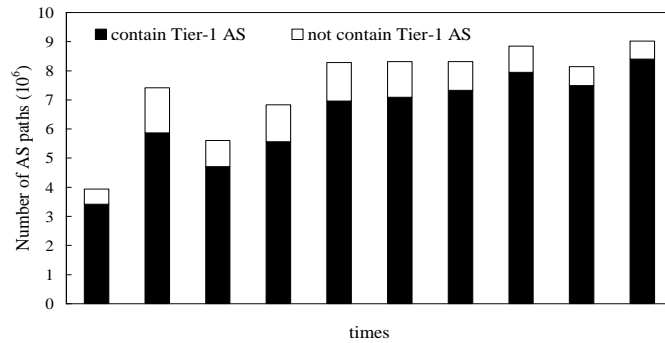


Figure 6. The fraction of AS paths which contain Tier-1 AS

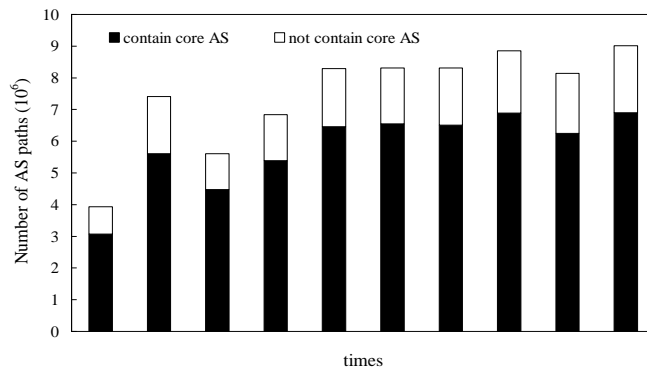


Figure 7. The fraction of AS paths which contain core AS

4. Evaluation of Prefix Hijacking Impact

In a prefix hijacking incident, attacker announces the IP prefix which belongs to victim network maliciously. This hijacking route propagates in the inter-domain routing system, too. The one who accepts it becomes an infector. Infectors redirect the data traffic aimed at victim to attacker. The more infectors exist, the more impact the hijacking attack has. Some infectors not only send their own traffic but also transmit others' traffic carried by them to the attacker. These infectors enlarge the impact of the hijacking. Through study of Hinge-Transmit property of BGP routing system we figure out which ASes they are and where they are located. In this section, we first define the quantitative measurement of BGP prefix hijacking impact; then evaluate the impact based on Hinge-Transmit property of BGP routing system; finally, validate the results by real prefix hijack incidents occurred in the Internet.

4.1. Prefix Hijacking Impact definition

When the attacker a wrongly announces victim v 's IP prefix, whether an AS i becomes infected depends on which route it selects to send traffic, the origin route $r_{i,v}$ or the hijacking route $r_{i,a}$. It follows the same process with normal routing decision without checking the origin AS whom the route is announced from. But we can figure the infected status by detecting which route AS i has chosen, like (6). If AS i chooses the hijacking route as the path that it send traffic through, AS i is infected, and its status function $Infect(a,i,v)$ is equal to 1; if AS i chooses the origin route, its status function is equal to 0 which means it's not infected.

$$Infect(a,i,v) = \begin{cases} 1, & \text{if } r_{i,v} = r_{i,a} \\ 0, & \text{if } r_{i,v} \neq r_{i,a} \end{cases} \quad (6)$$

In BGP routing table, if an AS becomes an infector, the origin AS of its AS path changes from victim to attacker, which is called a Multiple-Origin-AS conflict. For example, if AS500 announces the IP prefix 3.0.0.0/8 of AS400 maliciously, the infector's routing table changes from what is shown in Table 1 to that in Table 2. According to this, we calculate the number of infectors in a prefix hijacking incidents.

Table 2. BGP routing table entry after prefix hijack attack

Timestamp	Peer IP	Peer AS	Prefix	AS Path	...
1337920895	10.0.0.1	100	3.0.0.0/8	100 200 300 500	...

Furthermore, we can define the impact of a prefix hijacking attack $Impact(a,v)$ as the percentage of ASes who has become infectors in the AS set $ASSpace$ which consists of all ASes in the Internet, like (7).

$$Impact(a,v) = \frac{1}{|ASSpace|} \sum_{i \in ASSpace} Infect(a,i,v) \quad (7)$$

$Impact(a,v)$ is an appropriate measurement to describe the damage impact of an attack. The more infected ASes there are, the more data traffic will be wrongly redirected, the more impact a prefix hijacking incident will have.

4.2. Prefix Hijacking Impact Evaluation Method

A Prefix hijacking incident consists of an attacker, a victim and several infectors. This incident can be further divided into several scenes. Each of them contains one infector. A prefix hijacking scenario is shown in Fig. 8. This incident has 2 scenes which are Scene 1 consists of attacker, victim and infector 1; and Scene 2 consists of attacker, victim and infector 2.

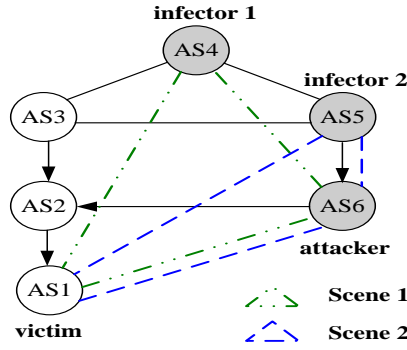


Figure 8. A prefix hijacking scenario

Impact of a prefix hijacking scene is presented like (8). From this function we can see that the impact is related to the Transmit factor of infector wrt. victim. The status function $Infect(a, i, v)$ is represented as engender factor of impact, while the Transmit factor $T_v(i)$ is enlarge factor. As shown in (9), the impact of an incident is larger or equal to the maximum impact of all the detected infected scenes. For simplicity, we consider this maximum value as the approximation of the evaluation result.

$$Impact(a, i, v) = T_v(i) \times Infect(a, i, v) \quad (8)$$

$$Impact(a, v) \geq \max_{i \in InfectedSpace} \{Impact(a, i, v)\} \quad (9)$$

The details of evaluation steps is listed as follows:

1. Analyze the topology locations of attacker and victim. By simulating prefix hijacking process which is similar to regular routing process, we can get all the possible infectors in a hijacking incident.
2. Calculate the impact of all scenes using Function (8).
3. Calculate the maximum value as the impact of the incident.

Based on the Hinge-Transmit property of BGP routing system, the Transmit factor of core AS wrt. its destination is very high, which means most of the traffic targets to the destination are transmitted by its core ASes. If one's core ASes are deceived by the hijacking route, there will be a lot of traffic supposed to be sent to the destination redirecting to the attacker instead. This enlarges the impact of a prefix hijacking attack.

4.3. Prefix Hijacking Impact Evaluation Results and Validation

There were some real prefix hijacking incidents occurred in the Internet. Their impacts were recorded by the routing tables of Route Views. We choose 10 incidents randomly; calculate their impacts based on our evaluation method; then validate them by Route Views data. Table 3 shows details of our evaluation results. And Fig. 9 presents validation result.

Table 3. Prefix hijacking impact evaluation results

id	attacker	victim	core AS set of victim	Is core AS infected?	impact
1	AS27506	AS20282	{AS3356, AS1239, AS174, AS701}	No	0.011
2	AS27506	AS33584	{AS3356, AS2914, AS701}	Yes	0.563
3	AS27506	AS7169	{AS174, AS3356}	No	0.062
4	AS27506	AS33477	{AS1299, AS2914, AS1239, AS3356, AS7018, AS2828}	No	0.101
5	AS9121	AS19198	{AS701}	Yes	0.404
6	AS9121	AS30576	{AS3356}	No	0.101
7	AS9121	AS31846	{AS174}	No	0.117
8	AS9121	AS30594	{AS3356, AS701}	No	0.035
9	AS9121	AS668	{AS701}	No	0.130
10	AS9121	AS19281	{AS701, AS2914, AS6461}	No	0.034




Figure 9. Prefix hijacking impact validation result

It's obvious that the impact of Incident 2 and 5 is much larger than others. In Incident 2, AS27506 wrongly announced the IP prefix 204.13.72.0/24 which belongs to AS33584 into the global routing system. Two of AS33584's core ASes which are AS701 and AS2914 chose to believe the hijacking route announced by attacker AS27506. This hijacking route propagated wildly in the whole network after that, launching an attack with impact as high as 0.563. In Incident 5, the IP prefix 65.164.53.0/24 which belongs to AS19198 was hijacked by AS9121. The core AS of victim was infected in this incident, which caused another impactive prefix hijacking attack. The other incidents in our experiment didn't cause the core ASes to be infected, so the damage in data reachability of these attacks is less than the previous two.

5. Conclusion

In this paper, we study BGP routing process under the control of routing polices. By evaluating a series of Transmit factors of AS which reflects the fraction of data traffic transmitted by the AS, we realize that BGP routing system has a Hinge-Transmit property. Based on this property, we evaluate the impact of BGP prefix hijacking attack, and find the root of matter why the impact differs a lot in different incidents. All of our findings and results are verified by the statistic data of BGP routing tables collected by Route Views project.

It is very important to realize the Hinge-Transmit property of BGP routing system. The few hinge ASes which transmit a large fraction of traffic are the keys of inter-domain routing system security. If these ASes are infected in BGP prefix hijacking, a lot of data traffic aimed at the victim network will be redirected to the attacker. To improve the security of inter-domain routing system, it is crucial and effective to protect these hinge ASes from misconfiguration and malicious attacks.

6. References

- [1] S. Kent, C. Lynn, J. Mikkelsen, and K. Seo, "Secure Border Gateway Protocol (S-BGP)", In Proceedings of Network and Distributed System Security Symposium (NDSS 2000), San Diego, California, 2000.
- [2] J. Ng, "Extensions to BGP to support secure origin BGP (soBGP)", in Internet Draft, Apr. 2004.
- [3] Z. Zheng, Z. Ying, and Y. C. Hu, "iSPY: Detecting IP Prefix Hijacking on My Own", In Proceedings of ACM SIGCOMM 2008 conference on Data communication, 2008, New York, NY, USA, 2008.
- [4] 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 ISOC NDSS'03, San Diego, CA, USA, Feb. 2003.
- [5] W. Xu and J. Rexford, "MIRO: multi-path interdomain routing", Computer Communication Review, vol. 36, pp. 171-182, 2006.
- [6] H. Ballani, P. Francis, and X. Zhang, "A study of prefix hijacking and interception in the Internet", ACM SIGCOMM Computer Communication Review, vol. 37, pp. 265-276, 2007.
- [7] M. Lad, R. Oliveira, B. Zhang, and L. Zhang, "Understanding resiliency of Internet topology against prefix hijack attacks", In Proceedings of the 37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, Edinburgh, Scotland, 2007.
- [8] W. Feng and G. Lixin, "On inferring and characterizing Internet routing policies", JOURNAL OF COMMUNICATIONS AND NETWORKS, vol. 9, pp. 350-355, 2007.
- [9] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, K. claffy, and G. Riley, "AS Relationships: Inference and Validation", <http://arxiv.org/abs/cs.NI/0604017>, 2005.
- [10] "CAIDA AS Relationships", <http://www.caida.org/data/active/as-relationships/>, 2006.
- [11] X. Dimitropoulos, D. Krioukov, B. Huffaker, K. Claffy, and G. Riley, "Inferring AS relationships: Dead end or lively beginning?", In Proceedings of 4th International Workshop on Experimental and Efficient Algorithms, Santorini Isl, GREECE, 2005.
- [12] "Tier 1 network - Wikipedia entry", http://en.wikipedia.org/wiki/Tier_1_network, 2006.
- [13] "Route Views Project Page", www.route-views.org, 2006.
- [14] B. Zhang, R. Liu, D. Massey, and L. Zhang, "Collecting the Internet AS-level topology", COMPUTER COMMUNICATION REVIEW, vol. 35, pp. 53-61, 2005.