On the Characteristics of BGP Multiple Origin AS Conflicts

Kwan-Wu Chin
School of Electrical, Computer and Telecommunications Engineering
University of Wollongong
Northfields Avenue, NSW, Australia
kwanwu@uow.edu.au

***Abstract*— The Internet has thousands of autonomous systems (ASs), each advertising one or more prefixes it owns. According to RFC 1930, each prefix should originate from a single AS. However, routing updates involving multiple ASs advertising reachability to a given prefix, so called multiple origin AS conflict (MOAS), are becoming prevalent. To this end, this paper analyzes and quantifies MOAS conflicts observed in routing updates collected over 21 days in January 2007, and presents ten reasons that explain why they occur.**

I. Introduction

The Border Gateway Protocol (BGP)[9] is the de-facto inter-domain routing protocol used to link autonomous systems (ASs). Each AS has a consistent routing policy, and uses BGP to connect its routers to other ASs. These routers exchange reachability information with other routers, and populate their respective routing table with the best paths to various networks on the Internet. An example reachability information is, 66.85.252.0/22 → [2905 701 6305]. The term on the right denotes the path vector of a given prefix. Each AS that has connectivity to the prefix/network adds its unique AS number to the path vector before forwarding the reachability information to its neighboring ASs. Hence, from the path vector, a router can easily determine the list of ASs packets will have to traverse through in order to reach the network 66.85.252.0/22. Moreover, the router can determine the AS that owns a given prefix, so called origin AS, by accessing the last AS on the path vector. For the prefix 66.88.252.0/22, its owner is AS 6305.

In this paper, we are interested in prefixes that originate from two or more origin ASs, so called multiple origin AS (MOAS) prefixes. In other words, more than one ASs claiming to be the owner of a given prefix. Take for example the following two reachability information, 68.254.214.0/24 → [12682 3491 7132 32380] and 68.254.214.0/24 → [2905 701 12026]. These two reachability information imply that the prefix/network 68.254.214.0/24 is located in the ASs 32380 and 12026. Hence, they violate RFC 1930 [3], which recommends that a prefix originates from one AS only. Unfortunately, in practice, this is not true. Therefore, it is important that we analyze MOAS conflicts, and understand why they occur.

This paper is structured as follows. We first present related works in Section II and highlight how our work extends Zhao et al. [10]'s seminal work on MOAS conflicts. Then, in Section III, we present our research methodology followed by characteristics of MOAS conflicts detected over 21 days. In Section IV, we present causes of MOAS conflicts, followed by our conclusions and future works in Section V.

II. Related Works

This paper revisits Zhao et al. [10]'s 2001 work, and presents updated results on the prevalence and characteristics of MOAS conflicts. We repeated most of Zhao et al.'s experiments, and also study MOAS conflicts with origin ASs located in different countries. From these experiments, we compare the results obtained in 2001 and 2007, and provide additional reasons that lead to MOAS conflicts.

After Zhao et al.'s work, researchers have developed methods to quickly detect whether a MOAS conflict is due to prefix hijack. This problem is particular critical given that Ballani et al. [1] recently showed that hijacking a prefix without disrupting traffic flow is possible. Readers interested in these methods are referred to [4][11], and their references. This paper however is focused on quantifying MOAS conflicts, and analyzing scenarios that lead to their occurrences, which could be due to reasons other than prefix hijack.

Apart from the aforementioned works, Huston [6] runs a web-site that automatically generates daily reports containing statistics related to BGP, one of which is multiple origin prefixes, aka MOAS conflicts. However, the site does not include reasons for these conflicts.

III. Methodology and Results

To study MOAS conflicts, we use BGP route updates collected by RouteViews[1]. Specifically, we process BGP updates dumped by the Oregon Internet Exchange (IX) from 1-January-2007 to 21-January-2007; compiled into 1943 compressed snapshots totalling 756 MByte in size. We process each snapshot for prefixes and record their origin AS. Besides that, we record the lifetime of a given prefix. Lastly, we use the WHOIS service to obtain the owner's name of a given prefix and AS number.

A. General

The total number of prefixes observed in the dumped BGP updates is 239,747, with 2,945 prefixes classified as MOAS conflict. Figures 1 and 2 show the number and proportion of these prefixes over the 21 days. The first day recorded 757 prefixes, however this number drops significantly with each passing day. From Figure 2, we can see that the number of prefixes with conflicts only constitutes a small percentage of the total number of prefixes advertised each day. Interestingly, the number of conflicts is more than twice that of the total recorded in the year 2001 [10]. Apart from that, when we plot prefix lengths, see Figure 3, we find that the most popular prefixes are 24 bits in length or /24, followed by those with lengths /18 to /23. This is consistent with the results from [10], but we observe more prefixes with lengths /18 to /23.

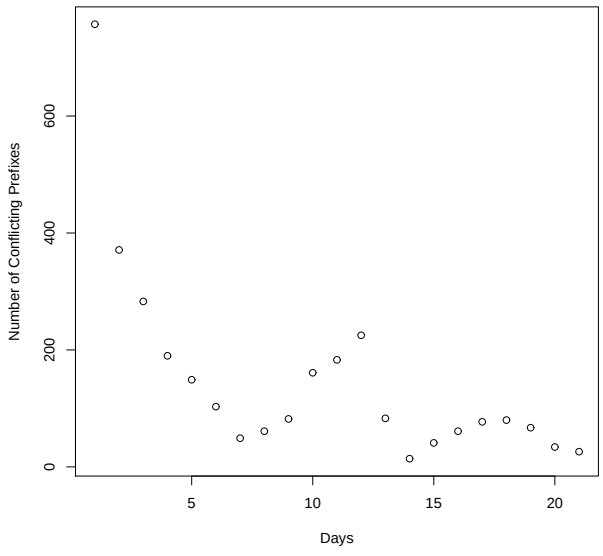

Fig. 1. Number of unique prefixes with MOAS conflicts over 21 days.

[1]http://www.routeviews.org

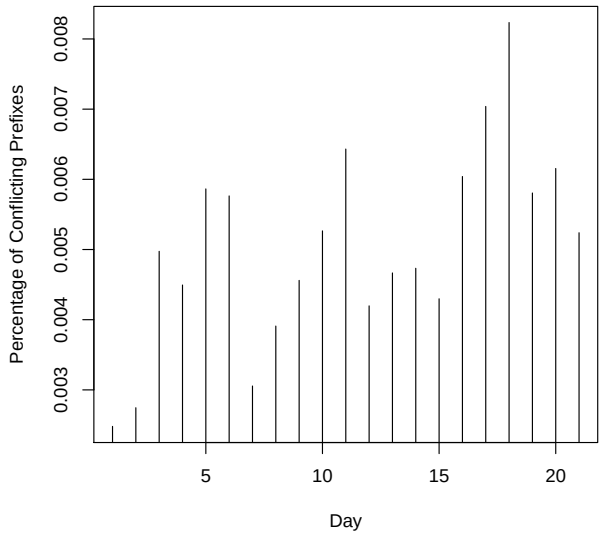

Fig. 2. Percentage of prefixes in conflicts.

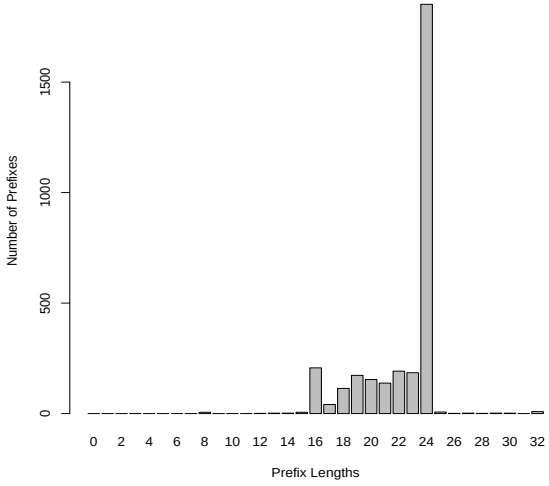

Fig. 3. Distribution of prefix lengths.

B. Distribution of Origin ASs

Figure 4 shows that most of the 2,945 prefixes are advertised by two ASs. Interestingly, there is a prefix advertised by 13 ASs! Upon closer inspection, this prefix turns out to be the 6to4 anycast relay router prefix, 192.88.99.0/24 [5]. We will elaborate further on other possible causes of MOAS conflicts in Section IV.

C. Path Characteristics

Another question of interest is the number of prefixes that have unique paths. In other words, those that do not share any common ASs. Figure 5 shows that the path vector of 1007 prefixes are unique. Also shown are the paths with at least one, two, etc ASs in common. Apart from that, 31 of them have one of the conflicting origin AS functioning as a transit AS. This usually occurs due

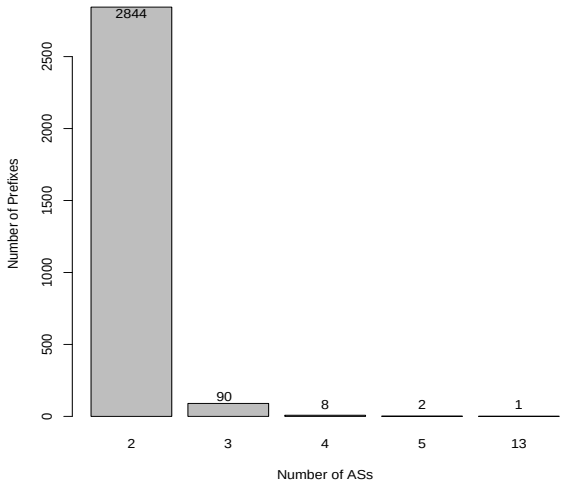

Fig. 4. Number of ASs advertising a prefix.

to traffic engineering purposes [10]. Note, we remove the AS which the collector is in before comparing paths for a given prefix. Otherwise, the number of paths with one common AS will be a lot higher since the same AS may have been used to collect route updates.

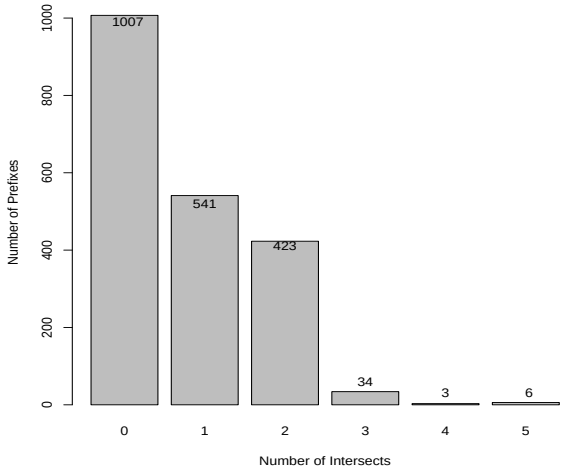

Fig. 5. Number of common ASs in paths.

D. ASs Countries

An interesting question is to determine the location of ASs advertising reachability to a given prefix. If both ASs are located in the same country, it is likely they have a peering arrangement over a link that does not run BGP. This is particularly true if both conflicting ASs are located in the same state or city. However, if they are located in different countries, then there is cause for concerns.

In total, there are 98 prefixes with origin ASs located in different countries. Figure 6 depicts the countries involved in advertising a prefix. For example, there are ASs in the US and China advertising the same prefix. There are also ASs in India and Bahrain originating the same prefix. From the figure, we can also quickly identify prefixes that warrant more investigations. For example, there is an edge between China (CN) and South Africa (SA). After further investigations, we found that an AS in China and South Africa is advertising ChinaNet's prefix 61.152.241.0, and it is unclear why a South African AS would be advertising the prefix. Lastly, notice that countries that are close by are well connected, especially those in Europe. Therefore, it is difficult to ascertain whether the detected MOAS conflicts from these countries are valid or invalid since the ASs involved may have a non-BGP link.

E. Conflicts Lifetime

Finally, we investigated the lifetime of prefixes. That is, we record the time when a prefix is first announced to the time it is withdrawn. Interestingly, all the detected MOAS conflicts have very short average lifetime. In fact, only 13.25 hours. This agrees with [10]'s observation, and it is unclear why MOAS conflicts are short-lived given that a majority of them are due to multi-homing. That is, they serve as backup or are used for traffic engineering purposes. Hence, logically they should be up for a longer period of time. This remains a future research question.

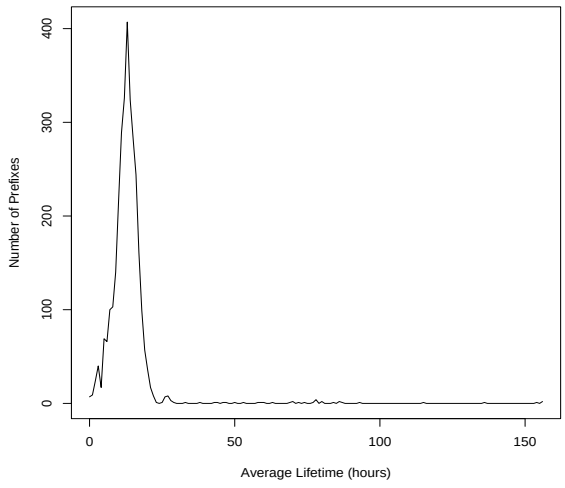

Fig. 7. Average Lifetime of Prefixes

IV. Explanations

The following sections outline key reasons that cause MOAS conflicts.

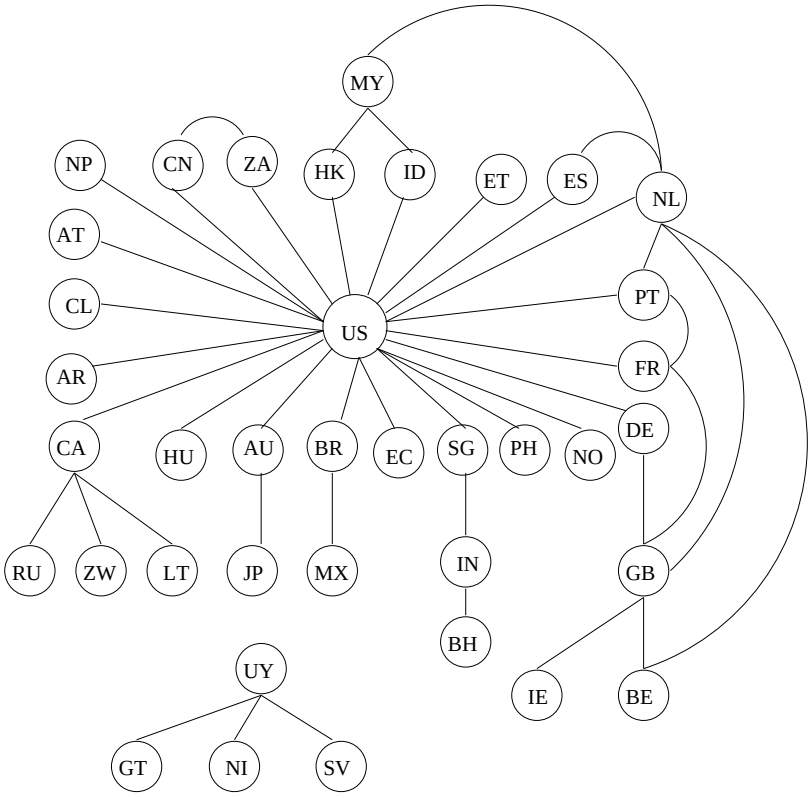

Fig. 6. Location of ASs causing MOAS conflicts. An edge indicates that at least one prefix have been advertised by an AS located in the countries connected by the edge. Note, the name corresponding to each country code can be resolved using http://www.iana.org/root-whois/index.html.

A. Anycasting

Anycasting is one of the many reasons that causes MOAS conflicts. The best example is the 6to4 anycast relay router prefix, 192.88.99.0/24 [5]. This prefix is advertised by various ASs to indicate they have a 6to4 relay router that can provide transit to IPv6 networks. Apart from that, we find that content distribution networks also employ anycasting. Examples include Akamai[2] and HighWinds[3], which we assume use anycasting to redirect traffic to their nearest networks/servers, thereby enabling fast and localized data delivery to their customers.

B. Internet Exchanges

Another cause of MOAS conflicts is due to ASs advertising IX prefixes. From our results in Section III-B, we found the two prefixes that originated from five ASs, 198.32.176.0 and 206.223.115.0, correspond to IX addresses. For example, 206.223.1165.0 belongs to Equinix Inc., which is a major IX provider. Similarly, for the eight prefixes that originated from four different ASs, three of them belong to IXs.

[2]http://www.akamai.net

[3]http://www.highwinds.com

C. Multi-National Companies

A multi-national company may advertise its prefix from countries it has offices in. For example, Glenayre Technologies has a subsidiary in China, called Glenayre Electronics, which uses the same prefix as its parent office in the US. To ensure reachability to all offices, ISPs in China and the US advertise a route to Glenayre Technologies' prefix. Having such configuration clearly reduces latency, since without this set up, traffic originating from China will be directed to the company's network in the US before being tunneled back to China.

D. AS Ownership

An organization may own multiple AS numbers, and chooses to advertise reachability to its prefixes from ASs it owns. For example, AT&T advertises the same prefix from three ASs it owns, namely 17228 17229 17227. Another example is Safelink Internet, which advertises the same prefix from two of its ASs; 33037 and 32444.

E. Multihoming

Multi-homing promotes reliability and also exposes a network to different pricing models offered by various tier-1 ISPs. For example, the Royal Bank of Canada,

which owns the prefix 192.234.99.0/24 and AS number 20069, is subscribed to both Sprint and AT&T. The bank peers with Sprint using BGP and has an IGP link with AT&T. Both Sprint and AT&T advertise the prefix 192.234.99.0/24. In a different example, an AS may advertise a more specific prefix to a secondary ISP to affect incoming traffic. We observe 220 occurrences of such practice.

F. Data Centers

A content provider, say Corp-A, may locate its servers at a company that specializes in providing high speed and reliable data service. For example, to take advantage of the services provided by the company Peer1[4], Corp-A places its content servers at one of Peer1's data centers. Peer1 then advertises reachability to company-A's servers, which bear company A's IP prefix. Hence, resulting in a MOAS conflict.

G. Satellite Networks

Developing countries using satellite links to connect to the Internet is another scenario that leads to MOAS conflicts. Figure 8 shows a set up consisting of two satellite links owned by AS-1 and AS-2 respectively, and are used by country-A to gain access to the Internet. To ensure connectivity to country-A, both ASs must advertise reachability to prefix-A.

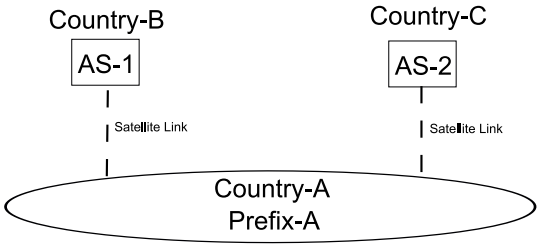

Fig. 8. A country connected to the Internet via satellite links.

In a similar scenario, consider a company that wants to keep its existing prefix, and does not run BGP. For reliability or economic reasons, the company subscribes to two ISPs, and advertise its prefixes to them using IGP [2]. To ensure reachability, both ISPs advertise the company's prefix along with their own prefixes, which result in MOAS conflicts. Certainly, there would be no conflict if the company has a valid AS number and uses BGP to peer with both ISPs.

[4]http://www.peer1.com

H. Umbrella Organization

In this scenario, a parent organization that owns a large chunk of IP addresses divides its address space to one or more child organizations; each may have a valid AS number. In turn, a child organization may decide to peer with another ISP that is different to the one providing service to the parent organization. In this scenario, the parent and child network will advertise the same address block, which result in a MOAS conflict. Moreover, as shown in Figure 9, the parent organization's export policies may allow a child organization's BGP updates to flow through un-filtered, resulting in both the parent and child AS advertising reachability to the same prefix.

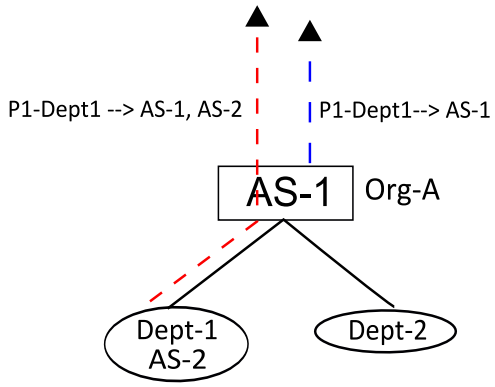

Fig. 9. An example of loose export policy of an umbrella organization.

I. Hijacks and Misconfigurations

MOAS conflicts are also caused by prefix hijacks [8][1] and mis-configurations [7]. In the former, a rogue AS may advertise a prefix belonging to another organization, in the hope of intercepting its traffic. Similarly, a mis-configuration causes unwanted incoming traffic. Apart from that, prefixes can be hijacked to send out spams [8]. In our study, we did not observe any evidence of hijacks nor any mis-configurations.

V. Conclusion

Prefixes with MOAS conflicts consistute a small, but growing, percentage of reachability information received by routers. In this paper, we have analyzed their characteristics, and discussed 10 reasons that explain their occurrences.

Currently, we are developing a tool to visualize MOAS conflicts. Apart from that, we are investigating reasons that explain the short lifetime of MOAS conflicts.

REFERENCES

[1] H. Ballani, P. Francis, and X. Zhang. A study of prefix hijacking and interception in the internet. In *Proceedings of ACM SIGCOMM*, Kyoto, Japan, Aug. 2007.

[2] H. Berkowitz, E.Davies, and L. Andersson. An experimental methodology for analysis of growth in the global routing table. Internet Draft: draft-berkowitz-tblgrow-00.txt, July 2001.

[3] J. Hawkinson and T. Bates. Guidelines for creation, selection, and registration of an autonomous system (as). RFC 1930, Mar. 1996.

[4] X. Hu and Z. M. Mao. Accurate real-time identification of IP prefix hijacking. In *Procs. of IEEE Security and Privacy*, Oakland, USA, 2007.

[5] C. Huitema. An anycast prefix for 6to4 relay routers. RFC 3068, June 2001.

[6] G. Huston. Bgp reports. Web-site: http://bgp.potaroo.net/index-bgp.html, 2007.

[7] R. Mahajan and D. Wetherall. Understanding BGP misconfigurations. In *Proceedings of ACM SIGCOMM*, Aug. 2002.

[8] A. Ramachandran and N. Feamster. Understanding the network-level behaviour of spammers. In *Proceedings of ACM SIGCOMM*, Jan. 2006.

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

[10] 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 *Procs. of ACM SIGCOMM Internet Measurement Workshop*, San Francisco, USA, Nov. 2001.

[11] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. F. Wu, and L. Zhang. Detection of invalid routing announcement in the internet. In *IEEE InternationalConference on Dependable Systems and Networks (DSN) 2002*, pages 59–68, Washington DC, USA, June 2002.