

University of Memphis

University of Memphis Digital Commons

Electronic Theses and Dissertations

11-26-2011

Understading Multiple Origin AS Conflicts

Yaoqing Liu

Follow this and additional works at: <https://digitalcommons.memphis.edu/etd>

Recommended Citation

Liu, Yaoqing, "Understading Multiple Origin AS Conflicts" (2011). *Electronic Theses and Dissertations*. 358.

<https://digitalcommons.memphis.edu/etd/358>

This Thesis is brought to you for free and open access by University of Memphis Digital Commons. It has been accepted for inclusion in Electronic Theses and Dissertations by an authorized administrator of University of Memphis Digital Commons. For more information, please contact khhgerty@memphis.edu.

UNDERSTANDING MULTIPLE ORIGIN AS CONFLICTS

by

Yaoqing Liu

A Thesis

Submitted in Partial Fulfillment of the

Requirements for the Degree of

Master of Science

Major: Computer Science

The University of Memphis

December 2011

To My Beloved Parents, Wife and Son.

Abstract

Liu, Yaoqing. M.S.. The University of Memphis. December 2011. Understanding Multiple Origin AS Conflicts. Major Professor: Dr. Lan Wang.

Internet routing problems are often difficult to detect and diagnose because one address prefix can be originated by multiple ASes. There is, however, no comprehensive analysis on the causes of Multiple Origin AS (MOAS) conflicts. In this paper, we study the characteristics of MOAS conflicts and compare them with those from 10 years ago. We also provide an in-depth examination of four MOAS causes – IXP, anycast, false origin AS, and origin-AS transition. Furthermore, we propose two heuristics to identify MOAS conflicts caused by false origin ASes and origin-AS transitions. The findings from our study and proposed heuristics can help us design effective mechanisms to distinguish legitimate MOAS conflicts from illegitimate ones, thus improving the reliability and security of Internet routing.

Table of Contents

1	Introduction	1
2	Background	2
3	Related Work	4
4	Data Sources and Preprocessing	5
5	Terminology	7
	(1) Definitions	8
	(2) Example	8
6	MOAS Characteristics	9
	(1) Number of MOAS Prefixes	9
	(2) MOAS Conflict Duration	11
	(3) Number of Origin ASes	13
	(4) Comparison with Previous Results	13
7	Identifying MOAS Causes	14
	(1) Internet Exchange Point Prefixes	15
	(2) Anycast Prefixes	18
	(3) Suspected False Origin AS Prefixes	20
	(4) Origin-AS Transition	23
	(5) Multi-homing	26
8	Conclusion	26
	Bibliography	30

List of Tables

1	RouteViews Collectors in Our Dataset	6
2	Number of Prefixes before and after Removing Reserved Prefixes . . .	7
3	Distribution of MaxOrigin for MOAS Prefixes	13
4	Comparison with Previous Results	14
5	IXP Prefix Statistics	16
6	IPv6-to-IPv4 Relay Anycast Prefix Statistics	19
7	Percentage of MOAS Prefixes with Suspected False Origin ASes . . .	22
8	The Summarized Results of MOAS causes	28

List of Figures

1	Valid MOAS scenario	3
2	Invalid MOAS scenario	4
3	Timeline of A MOAS Prefix with Three Origin ASes	9
4	Daily Count of MOAS Prefixes	10
5	Daily Count of New and Disappearing MOAS Prefixes	11
6	Cumulative Distribution of MOAS Prefix Conflict and Appearing Days	12
7	Histograms of MOAS Conflict Days	12
8	Persistence of IXP Prefixes Involved in MOAS Conflicts	17
9	MaxOrigin Distribution of IXP Prefixes Involved in MOAS Conflicts .	17
10	Persistence of Anycast DNS Prefixes Involved in MOAS Conflicts . .	20
11	Daily Count of MOAS Prefixes with Suspected False Origin ASes . .	22
12	Cumulative Distribution of Transition Periods (≤ 90 Days)	25
13	Comparison of Numbers of MOAS Conflicts	25
14	Percentage of Transition-Period MOAS Conflicts	26
15	Percentage of New Transition-Period MOAS Conflicts	27

1 Introduction

Internet routing problems are often difficult to detect and diagnose. One of the difficulties lies in the fact that one address prefix may be legitimately originated by multiple Autonomous Systems (ASes), i.e., the prefix may be involved in a legitimate MOAS (Multiple Origin AS) conflict [20]. Multi-homing, anycast and other common practices can lead to such legitimate MOAS conflicts. However, mis-configurations and attacks can also cause MOAS conflicts, which may lead to black-holes and traffic hijacking. Unfortunately, BGP has no built-in mechanisms to distinguish illegitimate MOAS conflicts from legitimate ones. In order to improve the reliability and security of Internet routing, we must first have a thorough understanding of the characteristics of different MOAS conflicts.

To this end, we conducted a comprehensive study of MOAS conflicts using four years of BGP data from Jan. 1, 2007 to Dec. 31, 2010. This data was collected from 193 distinct peer routers in 91 ASes by six Routeviews [12] collectors located in North America, Asia, Europe, and Africa. First, we obtained various statistics on MOAS prefixes from the data. We then compared these statistics with those from Zhao et al.'s study conducted ten years ago [20]. Moreover, we examined MOAS conflicts caused by Internet exchange points and anycast using the BGP data, Internet Exchange Point (IXP) address prefixes from PCH [13] and inferred anycast DNS addresses from Yu et al. [19]. Furthermore, we used two heuristics to identify MOAS conflicts caused by false origin ASes and origin-AS transitions. Finally, after identifying each type of MOAS conflicts from our data, we studied their characteristics such as the conflict duration and number of origins involved in the conflicts.

We have the following major findings: (1) there are more MOAS events occurring on a daily basis today compared to 10 years ago; (2) origin-AS transition is a major contributor to the large spikes in the observed MOAS conflicts; (3) IXP and anycast are unlikely to be major causes of the observed MOAS conflicts; and (5) 39.9% of

the one-day MOAS conflicts were caused by the false origin ASes identified using our proposed heuristic.

2 Background

An Autonomous System is a set of routers that share a single routing policy, and it is generally controlled by an Internet Service Provider (ISP). It uniquely identifies the network of the Internet, but the network does not have to be located at the same physical area. To the outside world, it is viewed as a single entity and each has an identity number assigned by Internet Registry. The number is called ASN number or AS number. Before 2008, the AS number was represented by 16-bit integers from 0-65535, and AS numbers 64512 through 65534 are used for private purposes or we call them private AS numbers. For example, the University of Memphis was assigned AS14048 as the AS number. Due to the increasing scale of the Internet, the 32-bit AS numbers were introduced. The AS23456 was allocated as transit AS for the handshake of 16-bit AS routers and 32-bit AS routers.

Border Gateway Protocol (BGP) is a path vector protocol used to carry routing information between two different ASes. The owner of each AS also have a set of Internet Protocol Prefix (IP prefix or prefix for short) assigned by the Internet Routing Registry (IRR). BGP uses a list of prefixes to announce reachability for each AS. In this scenario, the AS was called the origin of the prefixes. For example, the University of Memphis announces the prefix 141.225/16 with its own AS number AS14048 to the Internet, then we call AS14048 the origin of prefix 141.225/16 or AS14048 originates prefix 141.225/16. Note that the origin of one prefix derives from the BGP attribute-AS path and it may not be the real origin of the prefix, in other words, only if the origin runs BGP protocol with other ASes may we obtain the real origin of one prefix. However, in practice BGP is not running over every AS, and this fact leads to

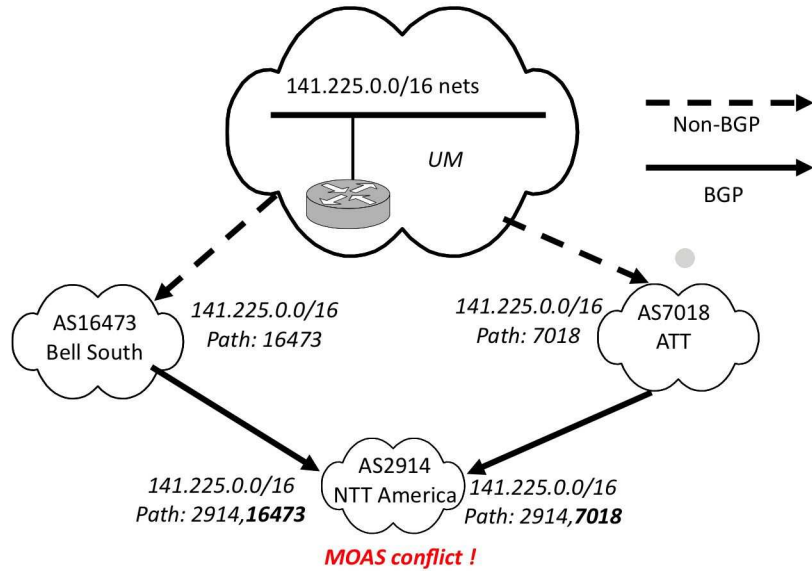


Figure 1: Valid MOAS scenario

inconsistent origin AS. One of the inconsistencies is the Multiple Origin AS Conflicts (MOAS), in which one prefix was announced by more than one AS, although RFC 1930 [4] recommends that one prefix should be originated from one AS. For example, as seen in Figure 1, the University of Memphis (UM) has two ISPs, ATT with AS7018 and Bell South with AS16473, and they announce the UM prefix 141.225/16 to NTT America with BGP announcement, there are no BGP sessions between UM and the two ISPs. In this case, MOAS conflict happens since NTT America will see two origins AS7018 and AS16473 for prefix 141.225/16. This case is valid in practice since there is no malicious behavior involved. However, it would be terrible that BGP may be hijacked by malicious users from a different origin AS, and it also would result in MOAS behavior. For example, as seen in Figure 2, UM was running BGP protocol with other ASes and announced the 141.225/16 to NTT America; Meanwhile, a malicious user hijacked the UT with AS3541, and announced the prefix 141.225/16 to NTT America as well through its ISP, then NTT America will see two origin ASes for prefix 141.225/16, this kind of practices, such as misconfiguration and hijacking,

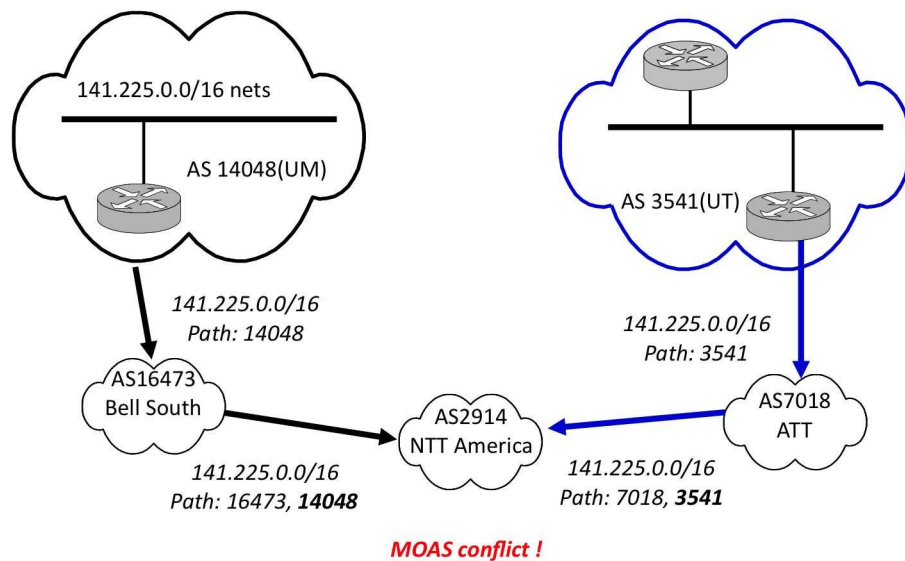


Figure 2: Invalid MOAS scenario

are invalid.

3 Related Work

The study most relevant to this work was conducted by Zhao et al. in 2001 [20]. They analyzed BGP MOAS conflicts from 11/08/1997 to 7/18/2001 and presented several metrics of the observed MOAS conflicts. They also classified the MOAS conflicts into three categories using relationships between the two AS paths in a conflict. Finally, they pointed out a number of possible causes of the MOAS conflicts, such as exchange point prefixes, anycast service, multi-homing without BGP or with private AS number, and misconfiguration. However, they did not further identify the contribution of each cause in their data.

Chin et al. [3] repeated the study in [20] using more recent data, and updated the statistics of MOAS conflicts that lasted for a very short time period. They extended some analysis on different countries in terms of origin ASes. They also observed that

a great number of short-lived prefixes exist, but did not identify their causes.

Mahajan et al. [10] studied two types of mis-configurations – Origin Misconfiguration and Export Misconfiguration for BGP route announcement. They mainly focused on studying very short-lived prefixes less than one day. Note that mis-configurations may or may not cause MOAS conflicts. In our work, we study MOAS conflicts caused by various legitimate practices, mis-configurations, and attacks.

Geoff Huston’s BGP Table Statistics website [7] provides a list of MOAS prefixes with their multiple origins and organization names. Hurricane Electric’s website [6] provides the Multi Origin Route Report. Team Cymru’s website [17] maintains a list of BGP inconsistent Origin ASNs and the corresponding prefixes with diagrams updated every two hours. They provide valuable information about the daily MOAS count and those ISPs involved. However, there is no comprehensive analysis what caused the MOAS and the contribution for each cause.

In Zhao et al. [21], they presented an approach to detect false route announcement resulting from invalid AS origination. They demonstrated that their approach can substantially mitigate the negative impact caused by invalid MOAS through simulation.

In another work Pei et al. [14] presented several causes of MOAS behavior, and one of them, a prefix temporarily multi-homed to both the old provider and a new provider, is the origin-AS transition considered in our paper.

4 Data Sources and Preprocessing

All of our results are based on the BGP data collected by the Routeviews project [12] that has been widely used for studying the global routing system. We use BGP RIBs and updates archived on 1,461 days between 01/01/2007 and 12/31/2010 from six collectors with a total of 193 peer routers. The detailed peer information for the six

collectors is shown in Table 1. From the table, we can see that the number of peer ASes for the six collectors ranges from 1 at kixp to 46 at roueview2.

BGP RIB tables are archived periodically at each collector every day. However, different collectors may use different archiving intervals and their clocks may slowly drift apart. In order to obtain the overall view of the global routing table, we need to aggregate BGP tables archived at the same time instance by the various collectors. Therefore, for each collector on each day, we downloaded one archived BGP table that is closest in time to midnight UTC and applied archived BGP updates to the downloaded table to produce the midnight table. We then aggregated the midnight tables from various collectors into one large table for each day. As a result, we obtained 1461 aggregated BGP tables. For example, for the rv2 collector on 9/1/2008, the BGP table closest to midnight UTC was timestamped 11:10pm. We applied all the BGP updates between 11:10pm on 9/1/2008 and 00:00am on 9/2/2008 to this BGP table to obtain the table at 00:00am on 9/2/2008 using BGP update rules for announcements and withdrawals.

Table 1: RouteViews Collectors in Our Dataset

Name	Location	Peer ASes	Peer Routers	BGP tables
rv2	U of Oregon, USA	46	68	1454
linx	London, GB	29	59	1459
isc	Palo Alto CA, USA	21	22	1461
eqix	Ashburn, VA, USA	19	30	1461
wide	Tokyo, Japan	10	13	1443
kixp	Nairobi, Kenya	1	1	1298
all	N/A	91	193	8576

We removed reserved address prefixes ([16,8,15,2,1]) from our data before further processing. These prefixes are not globally routable so they should not appear in the DFZ (Default Free Zone). However, some of the routeviews peer routers were not properly configured to filter these prefixes out in their peerings with the routeviews

Table 2: Number of Prefixes before and after Removing Reserved Prefixes

Date	IPv4 (before)	IPv4 (after)	IPv6 (before)	IPv6 (after)	Total (before)	Total (after)
01/01/2007	222813	222811	790	789	223603	223600
12/31/2010	356777	356640	4103	4103	360880	360743
2007-2010	713704	713306	6530	6529	720234	719835

collectors. This is likely a measurement artifact since the routers may be correctly configured with their operational peers, so we decided to remove the reserved prefixes. Table 2 shows the number of IPv4 and IPv6 prefixes before and after this cleaning procedure. There were a total of 720,234 prefixes before the cleaning and 719,835 prefixes left after the cleaning, i.e., 399 prefixes were removed all together during the four years. We also removed Private AS numbers (AS64512-AS65535) for similar reasons. Moreover, we removed AS23456 [18] from our data. This AS number is used in place of a 32-bit AS number for backward compatibility, when a router with a 32-bit AS number advertises its routes to a router that does not recognize 32-bit AS numbers. As such, AS23456 represents many different ASes. If we do not remove it, then this particular AS will appear to originate many prefixes even though it does not physically exist. Note that the removal of the above AS numbers has little effect on our results as their presence in the data is very limited.

For the purpose of identifying MOAS causes, we utilized the IXP prefix list from PCH [13] and a list of inferred anycast DNS addresses from the Internet Research Lab at UCLA [19].

5 Terminology

Before presenting our results, we introduce our terminology and use an example to demonstrate its usage.

(1) Definitions

MOAS prefix a prefix simultaneously originated by multiple ASes is called a *MOAS Prefix*. The prefix is said to be involved in a *MOAS conflict*. Note that a MOAS prefix may become a single-origin prefix from time to time.

New MOAS prefix on the first day when a prefix is simultaneously originated by multiple ASes, it is called a *New MOAS Prefix*. The prefix may have never appeared in the routing table previously or it may have existed as a single-origin prefix before it becomes an MOAS prefix.

Disappearing MOAS prefix on the last day when a prefix is simultaneously originated by multiple ASes, it is called a *Disappearing MOAS Prefix*. The prefix may continue to appear in the routing tables as a single-origin prefix or it may completely disappear from the routing table.

Prefix Appearing Days the total number of days when a prefix p appears in the routing table. This metric does not include the days on which p is absent from the routing table, e.g., when it is withdrawn by the origin AS(es).

MOAS Conflict Days the total number of days when a prefix p is simultaneously originated by multiple ASes. This metric does not include the days on which p is a single-origin AS.

(2) Example

Figure 3 shows when a particular prefix was originated by various ASes. It was originated by AS1 from t_0 to t_5 , AS2 from t_1 to t_2 , and AS3 from t_3 to t_4 . Therefore, it was a MOAS prefix from t_1 to t_2 and from t_3 to t_4 , and it was a single-origin prefix between t_2 and t_3 . It became a *new MOAS prefix* at t_1 and a *disappearing MOAS*




Figure 3: Timeline of A MOAS Prefix with Three Origin ASes

prefix at t_4 . Its number of *Prefix Appearing Days* is $(t_5 - t_0)$ and its number of *MOAS Conflict Days* is $(t_2 - t_1) + (t_4 - t_3)$.

6 MOAS Characteristics

In this section, we present an overview of the characteristics exhibited by the MOAS conflicts in our data.

(1) Number of MOAS Prefixes

We observed a total of 23,941 MOAS prefixes over the 4-year observation period, including 23,817 IPv4 prefixes and 124 IPv6 prefixes. The daily count of MOAS prefixes, however, is much smaller. As shown in Figure 4, the number of MOAS prefixes on each day ranges from 1,259 (01/30/2007) to 3,750 (09/11/2010) with a median of 1,905.

One can make a few more observations from Figure 4. First, there is an overall increasing trend in the daily count of MOAS prefixes. More specifically, the daily count gradually increased from $\sim 1,300$ in Jan. 2007 to $\sim 2,300$ in Dec. 2010 (there is a slight decrease towards the end of 2010). Second, the daily count has a small constant variation most of the time with a few large spikes.

In order to understand the variations in the observed MOAS count, we investigated the *new* and *disappearing* MOAS prefixes (see Section 5 for their definitions). Figure 5 shows that the number of new MOAS prefixes on each day ranges from 0 to 345 with




Figure 4: Daily Count of MOAS Prefixes

a median of 10, while that of disappearing MOAS prefixes ranges from 0 to 1,381 with a median of 9 (note that the y-axis in the top figure has a smaller range than that in the bottom figure). The low median values in the new and disappearing MOAS prefixes correspond to the small variations in the daily MOAS count. We believe that these small changes are mainly due to the false origination of address prefixes, which is usually short-lived and causes a small number of MOAS conflicts each day (see Section (3) for in-depth analysis of this type of MOAS conflicts). On the other hand, the large variations in the daily number of MOAS prefixes may have a different cause. We found that the large increases are typically caused by an elevated number of new MOAS prefixes over an extended period of time (from several days to several months), while the decreases are much more precipitous, often occurring in one day. This observation cannot be explained by false origination of address prefixes alone – although mis-configurations or attacks could cause large spikes in MOAS conflicts, these spikes tend to be very short-lived. After investigating four of the major spikes during our study period, we concluded that the main cause of these spikes is likely the simultaneous transition of many prefixes between their old and new origin ASes. Both origin ASes may announce the same prefixes for reliability, causing a large number

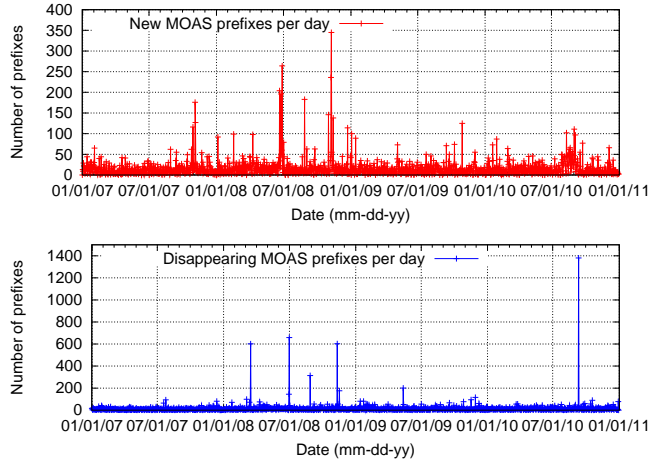


Figure 5: Daily Count of New and Disappearing MOAS Prefixes

of MOAS conflicts. If the transition takes many days, we will observe a spike over an extended period of time. Section (4) provides more supporting evidences for the above conclusion.

(2) MOAS Conflict Duration

Figure 6 shows the cumulative distributions of MOAS conflict days (top curve) and prefix appearing days (bottom curve) of the MOAS prefixes. From the bottom curve, we can see that 27.5% of the MOAS prefixes have Prefix Appearing Days for 1461 days, i.e., these prefixes were in the global routing system over the entire 4-year period. Ninety percent of the MOAS prefixes lasted more than 300 days, and only 0.3% lasted less than or equal than 15 days. These results suggest that most of the MOAS prefixes are very stable. From the top curve, one can tell that 75% of the MOAS prefixes experienced MOAS conflicts for less than 100 days and 50% experienced conflicts for less than or equal to 15 days. This is in sharp contrast to the bottom curve, which indicates that the majority of the MOAS conflicts did not last long, even though the prefixes themselves are stable.

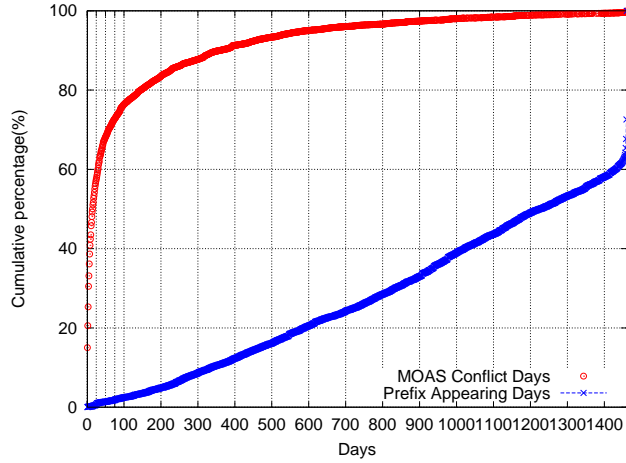



Figure 6: Cumulative Distribution of MOAS Prefix Conflict and Appearing Days



(a) Number of Days ≤ 150 days

(b) Number of Days ≤ 15 days

Figure 7: Histograms of MOAS Conflict Days

We take a further look at the short-lived MOAS conflicts in Figure 7(a) and 7(b). There are 19,231, 11,979, and 3,597 MOAS prefixes whose conflicts lasted less than or equal to 150 days, 15 days and 1 day, respectively. They represent 80%, 50%, and 15% of all the MOAS prefixes, respectively. In Section 7, we will examine the causes of these short-lived MOAS conflicts.

Table 3: Distribution of MaxOrigin for MOAS Prefixes

MaxOrigin	2	3	4	5	6	7	>7
MOAS Prefixes	23340	552	37	3	5	1	3

(3) Number of Origin ASes

We define MaxOrigin for an address prefix as the maximum number of ASes that simultaneously originate the prefix. Table 3 summarizes the distribution of MaxOrigin. We highlight the following results: (a) three MOAS prefixes, 192.88.99.0/24, 2001::/32, and 2002::/16, have a MaxOrigin more than seven. They are anycast prefixes used for IPv6-to-IPv4 relay service (Section (2) provides a more detailed analysis of their characteristics); (b) there is one prefix, 206.223.115.0/24, with a MaxOrigin of 7. This prefix belongs to Equinix (an Internet Exchange Point); and (c) 23,340 prefixes have a MaxOrigin of 2, which suggests that most MOAS conflicts involve only two origins.

(4) Comparison with Previous Results

Table 4 compares the characteristics of the MOAS prefixes in our data set with those in Zhao et al.’s work [20]. Our observation period was 1,461 days from 01/01/2007 to 12/31/2010. It is longer than Zhao’s observation period – 1,279 days from 11/08/1997 to 07/18/2001. We used 193 peer routers within 91 ASes from six global collectors, while they used 54 peer routers within 43 ASes from one collector. One observation is that our data has fewer MOAS conflicts (23,941 vs. 38,225) and a much lower percentage of one-day MOAS conflicts (15% vs. 36%). However, 11,358 MOAS conflicts in their study were due to one configuration fault, while there are no similar events in our data. If we take this event into account, the numbers in the two studies are much closer (23,941 vs. 26,867 and 15% vs. 9%). One can also observe that the

Table 4: Comparison with Previous Results

Metrics	Previous Data	Current Data
Observation period	1279	1461
Peer routers	54	193
Peer ASes	43	91
Median daily conflicts	683 (in 1998), 810.5 (in 1999), 951 (in 2000), 1294 (in 2001)	1362 (in 2007), 1589 (in 2008), 2019 (in 2009), 2346 (in 2010)
Expectation of conflict appearing days	30.9 (>0), 47.7 (>1), 107.5 (>9), 175.3 (>29), 281.8 (>89)	113.6 (>0), 133.5 (>1), 194.6 (>9), 271.3 (>29), 402.2 (>89)
Total conflicts	38225 (26867 ¹)	23941
One-day MOAS Conflict	36% (9% ²)	15%
Overall trend	increase	increase

¹ The total number of MOAS conflicts after removing those caused by a configuration fault

² The percentage of one-day MOAS conflict after removing those caused by a configuration fault

daily number of MOAS conflicts had an increasing trend in both their data and our data. Moreover, the median number of daily MOAS conflicts is much higher in our data. The above results suggest that there are more stable MOAS conflicts occurring on a daily basis in the routing system today.

7 Identifying MOAS Causes

In this section, we examine five potential causes of MOAS conflicts, including Internet Exchange Points, anycast, false origin AS, origin-AS transition, and multi-homing.

(1) Internet Exchange Point Prefixes

Zhao et al. [20] speculated that one cause of MOAS conflicts is Internet Exchange Points (IXPs), which provide Internet traffic exchange service between Internet Service Providers (ISPs). Because ISPs participating at an IXP typically share a common address prefix allocated to the IXP, the IXP prefix could be originated by any of the participating ISPs. However, we were surprised to learn from operators that IXP prefixes are generally not announced to the global routing table [9]. In fact, some IXPs have rules prohibiting the announcement of IXP address prefixes [9]. There is, however, a common exception – an ISP may announce an IXP address prefix to its customers for testing and diagnosis purposes. In this case, the prefix may be accidentally leaked to the global routing table, if the ISP or its customers do not filter their routing announcements properly; when multiple ISPs do so, there will be a MOAS conflict.

We obtained a list of 370 IXP prefixes from the website of Packet Clearing House (PCH) [13]. While this list may be incomplete for some countries (e.g., China), it is the most comprehensive public data on IXP that we could find. Moreover, given the total number of IXPs is inherently small, we consider this list as a good sample to gauge the MOAS behavior of IXP prefixes.

IXP Prefixes in Global Routing Table

Our results are summarized in Table 5. After removing three private address prefixes, we obtained 367 IXP prefixes (296 IPv4 and 71 IPv6) from the PCH website¹. Among these prefixes, we observed 137 prefixes (130 IPv4 and 7 IPv6) announced by the RouteViews peers over the four-year data period. Moreover, 38 (37 IPv4 and 1 IPv6) out of the 137 prefixes were involved in MOAS conflicts. In other words, *only 10.3%*

¹It is very rare for IXPs to use private IP address space in RFC1918 [15], because it will break traceroute and other diagnostic tools [9]. However, Maputo and Santiago de Compostela use the private address space for their IXPs [13].

of the studied IXP prefixes were involved in MOAS conflicts, and they represent only 0.16% of all the MOAS prefixes that we observed in our data. Finally, on the last day of our data period, we observed only 53 IXP prefixes (48 IPv4 and 5 IPv6) announced by the RouteViews peers, and 20 prefixes (19 IPv4 and 1 IPv6) were involved in MOAS conflicts. In summary, **announcement of IXP prefixes is not a pervasive behavior and most likely is *not* a leading contributor to MOAS conflicts.**

Table 5: IXP Prefix Statistics

Type	Total Pre-fixes	Announced	MOAS	Announced Last Day	MOAS Last Day
IPv4	296	130	37	48	19
IPv6	71	7	1	5	1
All	367	137	38	53	20

Characteristics of MOAS IXP Prefixes

Figure 8 shows two cumulative distributions of the 38 IXP prefixes that were involved in MOAS conflicts. The bottom curve is the cumulative distribution of the number of days on which a prefix was announced by the RouteViews peers (*Prefix Appearing Days*). One can find that (a) all the MOAS IXP prefixes appeared for more than 247 days; (b) more than 50% of them appeared for more than 1,000 days; and (c) nine prefixes were present in the global routing tables all the time (1,461 days). These results indicate that *such prefixes tend to stay in the global routing table for months*. The top curve is the cumulative distribution of the number of days on which a prefix is involved in MOAS conflicts (*MOAS Conflict Days*) (note that a MOAS prefix may be originated by a single AS on some days). It shows that over 50% of the MOAS IXP prefixes were involved in MOAS conflicts for more than 228 days. One prefix owned by Equinix Palo Alto was involved in MOAS conflicts over the entire 1,461-day

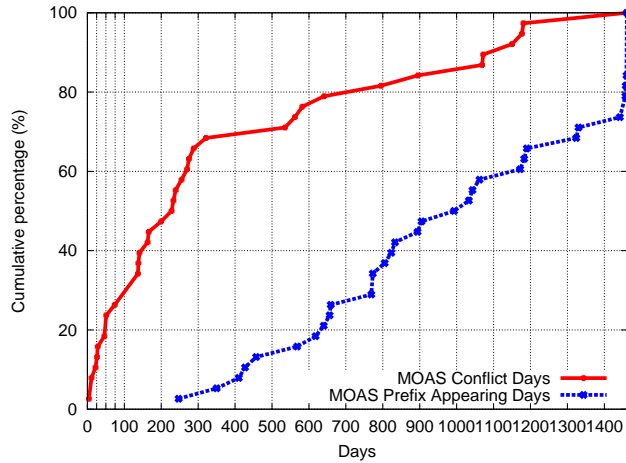


Figure 8: Persistence of IXP Prefixes Involved in MOAS Conflicts

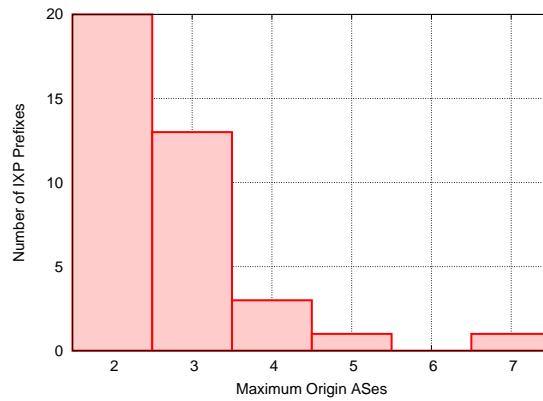


Figure 9: MaxOrigin Distribution of IXP Prefixes Involved in MOAS Conflicts

period. These results suggest that *the MOAS behavior of IXP prefixes tends to be long-lasting*. Figure 9 shows the distribution of the MaxOrigin for the MOAS IXP prefixes. We can see that (a) 20 of the 38 prefixes (52.6%) had two origin ASes; (b) 13 of the 38 prefixes (34.2%) had three origin ASes; and (c) one prefix (owned by Equinix) had a MaxOrigin of seven.

(2) Anycast Prefixes

Anycast is a technique to provide the same service in different physical locations using a common IP address space. To provide the anycast service, one address prefix is usually originated from multiple locations through BGP. As such, users obtain their service from the servers that have the best paths determined by BGP. Since the request load is distributed among multiple servers and users usually reach the closest servers, the service has better performance, reliability, and security. Today, anycast is commonly used for DNS, IPv4-to-IPv6 transition, and content delivery.

An anycast service can use a common AS number for all the locations, which does not cause MOAS conflicts. In fact, one operator on the NANOG mailing list remarked that “the overwhelming majority of anycast is done from homogeneous origin AS” [22]. This claim, if true, would suggest that anycast is not a leading cause of MOAS conflicts. While it is difficult to verify this claim, our results do provide some support for it (see Section (2) on Anycast DNS Prefixes).

On the other hand, the deployment of anycast services can lead to MOAS conflicts when each location originates the address prefix using a different origin AS number. This practice makes it much easier to diagnose problems as network operators can identify the specific AS that is providing service to their users. Therefore, McPherson et al. recommend the use of unique origin AS number per node [11]. If this recommendation becomes widely adopted, the number of MOAS conflicts due to anycast will likely increase.

In the remainder of this section, we provide detailed analysis on the MOAS behavior of anycast prefixes that provide 6to4 relay and DNS services. It is part of our future work to identify anycast prefixes that provide other types of services and study their MOAS behavior.

6to4 Relay Anycast Prefixes

RFC3068 specified three anycast prefixes that provide IPv6-to-IPv4 relay service: 192.88.99.0/24, 2001::/32 and 2002::/16 [5]. We summarize their statistics in Table 6. Among the three prefixes, 192.88.99.0/24 has the most ASes originating it — it was simultaneously originated by at least 6 ASes and at most 21 ASes, and it has a total of 56 origin ASes. Being IPv6 prefixes, 2001::/32 and 2002::/16 have fewer origin ASes in all three categories.

Table 6: IPv6-to-IPv4 Relay Anycast Prefix Statistics

Prefix	Min ASes in Conflict	Max ASes in Conflict	Total Origins
192.88.99.0/24	6	21	56
2001::/32	2	10	26
2002::/16	4	13	40

Anycast DNS Prefixes

DNS service is a natural candidate for anycast, especially for those root and gtld DNS servers with extremely high load. Recently, Yu et al. used DNS query and traceroute to infer DNS anycast addresses [19]. We obtained their list of 343 anycast DNS addresses, which are distributed in 106 address prefixes. After matching these prefixes with our MOAS data, we found that 20 of the 106 prefixes were involved in MOAS conflicts (78 addresses are distributed in these 20 prefixes). In other words, only 18.9% of these anycast DNS prefixes have MOAS behavior. This is one evidence supporting the claim that most anycast services use homogeneous origin AS. Among the 20 MOAS prefixes, 19 prefixes had a MaxOrigin of two and 1 prefix had a MaxOrigin of six.

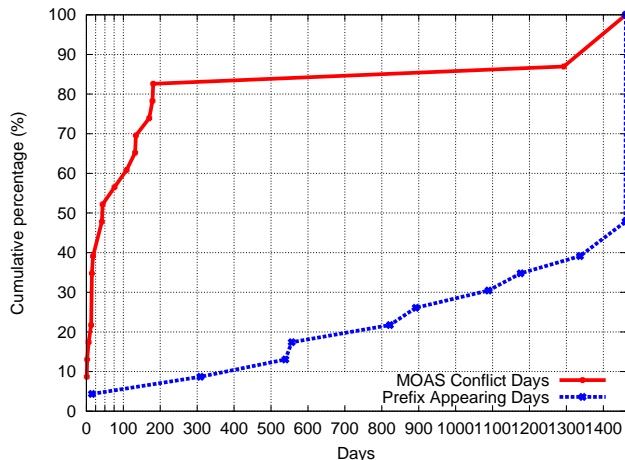


Figure 10: Persistence of Anycast DNS Prefixes Involved in MOAS Conflicts

(3) Suspected False Origin AS Prefixes

Previous studies ([20, 3]) discovered a large number of short-lived MOAS conflicts and found that many of them were caused by misconfigured BGP policies. However, it remains a challenge to automatically distinguish false origin ASes from legitimate ones. In this section, we propose a heuristic to identify false origin ASes in MOAS conflicts. The false origins may be caused by mis-configurations or attacks, but we do not attempt to infer the specific causes.

Heuristic

Routing policies on address prefix origination are usually long-term, so a legitimate origin AS should last for months. When an AS falsely originates a prefix, it is likely to withdraw the announcement after a short period of time, either to prevent the potential consequences (in case of misconfigurations) or to avoid detection (in case of attacks). Therefore, the duration of prefix origination by an AS can be one indicator for identifying false origin ASes. However, we cannot use this indicator alone as we may observe the same behavior with legitimate backup origin ASes. More specifically,

when the primary origin AS for an address prefix fails, some ISPs may adopt the backup origin AS while the others may still use the primary origin AS due to routing convergence delay, leading to a MOAS conflict. When the primary origin AS recovers, the backup origin AS disappears, not due to withdrawal of the backup origin AS, but due to BGP’s best path selection. In this case, we also observe a MOAS conflict and the quick disappearance of one of the origins. Fortunately, the two types of MOAS conflicts have slightly different behavior: a false origin AS usually appears only once, while a backup origin AS may appear from time to time, albeit for a short amount of time during each episode. We use this observation to differentiate them.

We label an origin AS as a *suspected false origin* for an address prefix if it satisfies three requirements: (a) the origin AS appeared on consecutive days, i.e., not intermittently. This rule is used to differentiate false origins from legitimate backup origins; (b) the origin AS appeared for a short period no more than a predefined number of days. After some experimentation, we decided to use seven days as the threshold (see Section (3)); and (c) the prefix had at least one other origin AS appearing earlier than this origin AS, as well as one appearing later than this origin AS.

Results

We found that 2,573 MOAS prefixes have one or more suspected false origin ASes and 2,239 MOAS prefixes become single-origin prefixes after removing the false origins (i.e., their MOAS behavior is solely caused by the false origins). Figure 11 shows the number of MOAS prefixes with suspected false origins on each day. This daily count ranges between 0 and 64 with a median of 2. These results suggest that false origin ASes were not a major cause of the MOAS conflicts in our data set. However, Table 7 shows that a considerable percentage of the *short-lived* MOAS conflicts were caused by them. For example, 39.9% of the one-day MOAS conflicts involved the false origin ASes that we identified.




Figure 11: Daily Count of MOAS Prefixes with Suspected False Origin ASes

Table 7: Percentage of MOAS Prefixes with Suspected False Origin ASes

Days	1	2	3	4	5	6	7
Total pre- fixes	3597	1325	1139	1240	631	722	599
False-origin prefixes	1435	351	255	170	117	106	147
Percentage	39.9	26.5	22.4	13.7	18.5	14.7	24.5

Note that due to the limitation of our methodology, we may miss those MOAS conflicts that appeared between consecutive midnight BGP snapshots and these short-lived MOAS conflicts are likely to be caused by false origins.

Validation

It is extremely challenging to validate the above results on false origin ASes due to the following reasons. First, our results involve a large number of ISPs. Second, ISP contact information in whois databases may be out of date. Third, operators may be reluctant to share sensitive information. Nevertheless, we are currently working on an automatic tool to obtain validation information from network operators for each suspected false origin, and we expect to get some preliminary data in the near future. Meanwhile, we used BGP routing tables from January 2011 to April 2011 to indirectly verify our results. More specifically, we checked if each suspected false origin AS reappeared in the new data. If an origin AS did reappear, it most likely is not a false origin and we call it a false positive. *We did not find any false positives in our results.* Although this is not a direct validation, it does strengthen our confidence in the results. We also used the new data to experiment with different time threshold for our heuristic, and found that any time threshold less than eight days has zero false positives. Therefore, seven days appears to be an appropriate time threshold for our heuristic method.

(4) Origin-AS Transition

An organization may decide to switch the origin AS of its address prefix from one to another. In order to maintain connectivity during the transition period, both the old and new origin AS may announce the same prefix for a period of time. We call this kind of MOAS behavior “Transition-Period MOAS Conflicts” and such kind of prefixes “Transition-Period Prefixes”.

Heuristic

To detect Transition-Period MOAS conflicts, we identify prefixes with two origin ASes that partially overlapped for a period of time. Moreover, the first origin should appear before the second one appeared and disappear before the second one disappeared. Furthermore, the first origin should not reappear later in the study period. Finally, the overlapping period should be short to filter out multi-homing cases. We chose 90 days as our threshold after some experimentation. For example, the prefix 8.22.224.0/21 was announced by AS47088 from 7/15/2009 to 9/19/2009, and then the origin switched to AS3361 from 9/19/2009 to 12/31/2010. This prefix satisfies all of our requirements, so our algorithm will consider its two origin ASes to be involved in a Transition-Period MOAS Conflict. In this example, the transition period is only one day (9/19/2009).

Results

We identified a total of 7,807 transition-period prefixes using our heuristic. Furthermore, we used BGP data from Jan. 2011 to Apr. 2011 to verify our results. If the old origin AS of an identified prefix reappeared in 2011, it is likely that the previous behavior was not origin-AS transition. We found the old origin ASes of 52 prefixes reappearing in 2011, which represent only 0.67% of the identified prefixes. This indicates that our heuristic is relatively robust.

Figure 12 shows the cumulative distribution of the duration of those transition-period conflicts that lasted no more than 90 days. 60% of transition-period conflicts lasted less than 10 days, 80% of them lasted less than 30 days, and 95% of them lasted less than 60 days. In other words, the transition period is typically no more than a few weeks.

Figure 13 shows the comparison between the daily count of overall MOAS conflicts (top blue curve) and Transition-Period MOAS conflicts (bottom red curve), which




Figure 12: Cumulative Distribution of Transition Periods (≤ 90 Days)

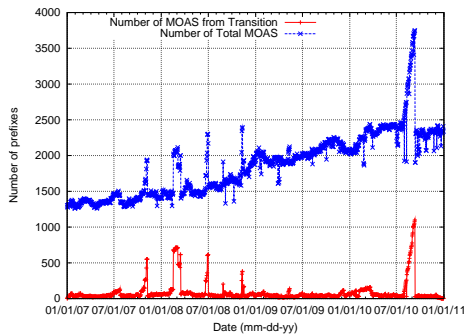


Figure 13: Comparison of Numbers of MOAS Conflicts

ranges from 0 to 1,220 with a median of 259. One interesting observation is that the large spikes in this figure correspond in time to the spikes in the daily count of overall MOAS prefixes. Figure 14 further shows that the Transition-Period MOAS prefixes typically constitute less than 5% of the overall MOAS prefixes, but this number reached 35% during those spikes. However, these two figures cannot yet lead to the conclusion that the increases during the spikes were mainly caused by the transition-period prefixes. Therefore, we investigated the new MOAS conflicts for each spike period. Figure 15 shows the percentage of the new transition-period MOAS conflicts over the total number of new MOAS conflicts for the four spike periods, 10/13/2007-11/8/2007, 2/11/2008-3/16/2008, 6/21/2008-7/1/2008 and 7/30/2010-9/11/2010. From the figures, we can see that the new transition-period conflicts constituted to up to 90% of the overall new MOAS conflicts with an average of more




Figure 14: Percentage of Transition-Period MOAS Conflicts


than 50%. Thus we can conclude that the origin-AS transition was the leading cause of the four sudden increases of MOAS daily conflicts.

(5) Multi-homing


After removing all the suspected false origins and those origins involved in transition-period MOAS conflicts, we have a total of 13,924 MOAS prefixes remaining, including 13,851 IPv4 prefixes and 73 IPv6 prefixes. Since our study of IXP and anycast prefixes led to the conclusion that they most likely are not major causes of MOAS conflicts, we suspect that the majority of the remaining MOAS prefixes are caused by multi-homing. When an organization has multiple providers, its address prefix may be originated by the providers if the organization does not own an AS number or it does not support BGP routing. This type of MOAS conflicts should be long-lasting, but we do not have a reliable method to distinguish them from those caused by anycast. It remains our future work to identify MOAS conflicts caused by multi-homing.

8 Conclusion


MOAS conflicts result from both valid and invalid causes, and this fact poses challenges to the Internet stability and security. The purpose of studying those with valid




(a) First Spike (10/13/2007-11/8/2007)



(b) Second Spike (2/11/2008-3/16/2008)



(c) Third Spike (6/21/2008-7/1/2008)



(d) Fourth Spike (7/30/2010-9/11/2010)

Figure 15: Percentage of New Transition-Period MOAS Conflicts

causes is to understand the BGP operational behavior. The analysis would help operators to view how BGP protocol performs and what really happened in the Internet for different services, such as IXP, Anycast and so forth. We study the invalid causes for the purpose of detecting malicious behaviors caused by BGP protocol deficiency, such as misconfiguration and hijacking. Moreover, sometimes we have to combine the two kinds of causes together to figure out the solution of a MOAS problem.

In this paper, we studied MOAS conflicts using four years of BGP data from Jan. 2007 to Dec. 2010. Compared with ten years ago, there are more stable MOAS conflicts each day, indicating that more organizations have adopted practices that lead to legitimate MOAS conflicts. We also found that the overall MOAS conflicts per day still has an increasing trend towards the end of our data. Moreover, we found that origin-AS transition, rather than false origin ASes, may be a major cause of the spikes in the observed daily MOAS conflicts. Furthermore, our analysis supports the conjecture that IXP and anycast are not major contributors to MOAS conflicts. Finally, we developed two heuristics to detect false origin ASes and origin-AS transitions. The summarized results are demonstrated in table 8.

Table 8: The Summarized Results of MOAS causes

Causes	IXP	DNS and 6to4 Relay Anycast	False Origin	Transition Period	Multi-homing and others
Number	38	23	2,573	7,807	13,924

To categorize specific MOAS prefixes into different classes is extremely challenging and our approaches have limitations. We have not found a very efficient way to detect anycast prefixes. We were using telnet service to login a few globally distributed Looking Glass Servers-BGP route Servers and issue traceroute commands to those anycast-like prefixes. However, this approach required the participation of each IP address for those prefixes, and thus it was slow and time-consuming. Moreover, the

inference results in the paper just had a preliminary verification and justification based on the data year 2011. They have not been verified by the ground truth from the operators. Therefore, most of time we underestimate the number of prefixes for each category.

The future work includes the following: (1) design a new heuristic method to identify MOAS conflicts caused by multi-homing; (2) validate these results using operator feedback; and (3) use these findings and heuristics to design mechanisms to detect illegitimate MOAS conflicts in real-time.

Bibliography

- [1] Z. Albanna, K. Almeroth, D. Meyer, and M. Schipper. RFC3171: IANA Guidelines for IPv4 Multicast Address Assignments, 2001. <http://tools.ietf.org/html/rfc3171>.
- [2] S. Bradner and J. McQuaid. RFC2544: Benchmarking Methodology for Network Interconnect Devices, 1999. <http://tools.ietf.org/html/rfc2544>.
- [3] K.-W. Chin. On the characteristics of BGP multiple origin AS (MOAS) conflicts. In *Proceedings of The Australasian Telecommunications Networks and Applications Conference (IEEE ATNAC'07)*, Dec. 2007.
- [4] J. Hawkinson and T. Bates. Guidelines for creation, selection, and registration of an Autonomous System (AS). *RFC 1930*, Mar. 1996.
- [5] C. Huitema. RFC3068: An Anycast Prefix for 6to4 Relay Routers, 2001. <http://www.ietf.org/rfc/rfc3068.txt>.
- [6] Hurricane Electric. Multi Origin Route Report. <http://bgp.he.net/report/multi-origin-routes>.
- [7] G. Huston. BGP Multi-Origin Prefixes. <http://bgp.potaroo.net/as6447/bgp-multi-org-prefix.txt>.
- [8] IANA. RFC1797: Class A Subnet Experiment, 1995. <http://tools.ietf.org/html/rfc1797>.

- [9] Y. Liu. Internet Exchange Point (IXP) questions, NANOG Mailing List, Feb. 2011. <http://www.gossamer-threads.com/lists/nanog/users/137930>.
- [10] R. Mahajan, D. Wetherall, and T. Anderson. Understanding BGP misconfiguration. In *Proceedings of the ACM SIGCOMM*, Aug. 2002.
- [11] D. McPherson, R. Donnelly, and F. Scalzo. Unique Per-Node Origin ASNs for Globally Anycasted Services. *Work in Progress, draft-ietf-grow-unique-origin-as-01.txt*, July 2011.
- [12] U. of Oregon. Routeviews Archive. <http://archive.routeviews.org>.
- [13] Packet Clearing House. PCH Website. <http://www.pch.net>.
- [14] D. Pei, W. Aiello, A. Gilbert, and P. McDaniel. Origin disturbances in bgp. Technical report, 2004. Technical Report TD-62TJJF8, ATT Labs - Research, Florham Park, NJ.
- [15] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear. RFC1918: Address Allocation for Private Internets, 1996. <http://tools.ietf.org/html/rfc1918>.
- [16] J. Reynolds and J. Postel. RFC1700: Assigned Numbers, 1994. <http://tools.ietf.org/html/rfc1700>.
- [17] Team Cymru. BGP Inconsistent Origin ASN List. http://www.cymru.com/BGP/incon_asn_list.html.
- [18] Q. Vohra and E. Chen. RFC4893: BGP Support for Four-octet AS Number Space, 2007. <http://tools.ietf.org/html/rfc4893>.
- [19] Y. Yu, J. Cai, E. Osterweil, and L. Zhang. Measuring the placement of dns servers in top-level-domain. Submitted to IMC 2011.

- [20] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. Wu, and L. Zhang. An analysis of BGP multiple origin AS (MOAS) conflicts. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop 2001*, Nov. 2001.
- [21] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. Wu, and L. Zhang. Detection of invalid routing announcements in the Internet. In *Proceedings of the International Conference on Dependable Systems and Networks*, June 2002.
- [22] Z. Zhu. IPv4 Anycast? NANOG Mailing List, Apr. 2009. url-
<http://seclists.org/nanog/2009/Apr/760>.