A Longitudinal Study of BGP MOAS Prefixes

Quentin Jacquemart, Guillaume Urvoy-Keller, Ernst Biersack

A Longitudinal Study of BGP MOAS Prefixes

Quentin Jacquemart[1], Guillaume Urvoy-Keller[2], and Ernst Biersack[1]

[1] Eurecom, Sophia Antipolis
[2] Univ. Nice Sophia Antipolis, CNRS, I3S, UMR 7271, 06900 Sophia Antipolis

Abstract. An IP prefix can be announced on the Internet from multiple endpoints, possibly leading to so-called MOAS (Multiple-Origin AS) prefixes. Long-lived MOASes are traditionally considered to be the result of network topology engineering such as prefix multihoming. Short-lived MOAS are commonly attributed to be the result of router misconfigurations.
In this article, we look at MOAS prefixes in the long term and seek the patterns behind these situations. We first revisit previous work by looking at the duration of MOAS events. We group these events according to the prefix announced and show that short-lived MOASes are not due to misconfigurations, but to origin instability or route flapping. We also identify topology patterns that result in MOAS prefixes and use them to classify these events. We show that, contrary to popular belief, multihoming is neither the main use case leading to MOAS, nor the most popular pattern. Finally, we look at the evolution of these observations by analysing data collected 10 years apart.

1 Introduction

The Internet is composed of a set of interconnected independent networks, known as **Autonomous Systems** (ASes), that exchange reachability information containing **IP prefixes** through the use of **BGP** (Border Gateway Protocol). A MOAS event, MOAS prefix, or simply **MOAS** (Multiple-Origin AS) is the result of an IP prefix p being announced simultaneously from multiple endpoints. Even though RFC1930 [1] discourages MOAS situations, Zhao et al. presented a series of legitimate network engineering practices that lead to MOAS prefixes in [2], such as prefix multihoming and the use of anycast. These are expected to create **long-lived** MOAS events. Zhao et al.'s analysis also uncovered a number of **short-lived** MOAS events whose root causes are unclear, and were attributed to router misconfigurations. In this paper, we provide an in-depth study of MOASes with which we clarify and quantify the root causes behind long-lived and short-lived MOAS events.

First, we use a two-sided approach. By considering MOAS events individually, we revisit previous work by focusing on MOAS durations. Then, we provide the first study of MOAS events as groups of events related to their prefixes. With this study, we show that short-lived MOASes are less numerous than previously reported in [2]: many short-lived MOAS events are actually repeated events related to a small set of prefixes due to instability or route flapping.

Second, we introduce a taxonomy of MOASes into distinct MOAS patterns – **peering**, **classical**, and **me-too** MOAS – and study their prevalence and temporal characteristics. With these patterns, we show that the majority of MOASes are fake MOASes

that are the result of loosely-defined, or outdated policies. The traditional MOAS shape only amounts to 30% of all cases.

Finally, we also look at the evolution of these findings by relying on the analysis of two full years of measurement collected 10 years apart, in 2002 and 2012 respectively, in order to underline discrepancies that may arise due to changes in standard practices, or due to the global evolution of the Internet.

2 Methodology and Dataset

2.1 Definitions

The Internet is composed of tens of thousands of interconnected independent **AS**es (Autonomous Systems). Inter-AS routing is done with BGP (Border Gateway Protocol), defined in RFC4271 [3]. BGP update messages are exchanged among BGP routers in order to propagate reachability information – containing IP prefixes and attributes – between ASes. One of these attributes is the **AS path**. When propagating a route, each router *prepends* its globally unique **ASN** (AS number) to the **AS path**. As a result, the *rightmost* ASN in the AS path is the **origin** of the route (unless the route was aggregated).

A **MOAS prefix** (Multiple-Origin AS) is the result of a prefix p being simultaneously originated from multiple ASes. In other words, at a given point in time, the AS paths for p end by a set $\mathcal{O}(p)$ of multiple origin ASNs, so that $\mathcal{O}(p) = \{a_1, \ldots, a_n\}$. For example, using Fig. 1, $\mathcal{O}_{]t_0,t_1[}(p) = \{1\}$, $\mathcal{O}_{]t_1,t_2[}(p) = \{1,2\}$, $\mathcal{O}_{]t_2,t_3[}(p) = \{2\}$, $\mathcal{O}_{]t_3,t_4[}(p) = \varnothing$, and so on. It is important to stress that MOASes only occur for the same prefix p. In particular, any prefix q more specific than p with a different origin than that of p is not defined as a MOAS prefix, but as a MOAS *sub*prefix (alternatively *sub*-MOAS), which we will not discuss in this document.

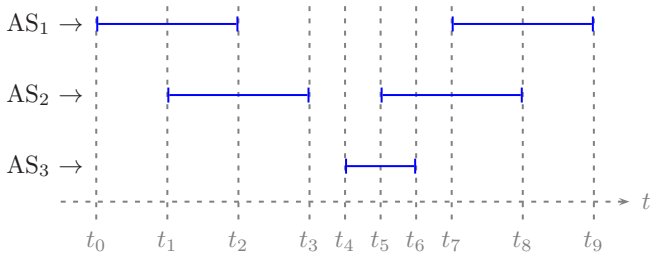

Fig. 1: Example of announcements for a prefix p

The literature defines **MOAS event duration** as the duration of a single MOAS event. In Fig. 1, the durations of the three MOAS events are $t_2 - t_1$, $t_6 - t_5$, and $t_8 - t_7$. MOASes are usually classified according to this metric with the following terminology [2]: **short-lived MOAS events** last less than 1 day, while **long-lived MOAS events** last more than 1 day.

If p is not a MOAS, but p is still present in the routing tables, p is a **SOAS** (Single-Origin AS), meaning that p is originated by a single AS. In Fig. 1, this happens during $]t_0, t_1[$, for example. If p is not included in the routing tables, we will say that p is **down** (Fig. 1 during $]t_3, t_4[$). This does not imply that traffic destined to p cannot be routed, because a set of covering prefixes could be used to forward the traffic. By contrast, a prefix is **up** whenever it is a SOAS or a MOAS.

We define the **lifetime** of a prefix p as the difference between the timestamp at which the prefix was last withdrawn (that is, the timestamp at which the prefix goes down for the last time) and the timestamp at which the prefix was first announced. The lifetime of p in Fig. 1 is simply $t_9 - t_0$. On the other hand, the **uptime** of p is defined as the total duration during which the prefix was advertised. In Fig. 1, the uptime of p is $t_3 - t_0 + t_9 - t_4$.

In Section 4.2, we introduce a new metric that we call the **MOAS duration per prefix** which is defined as the sum of the durations of the individual MOAS associated with this prefix. Using Fig. 1, the MOAS duration for prefix p is $t_2 - t_1 + t_6 - t_5 + t_8 - t_7$.

2.2 BGP Dataset

In order to study MOASes, we use data from RIPE RIS's [4] route collector located in Amsterdam (`rrc00`), which has above 40 geographically diverse peers. We retrieve the update messages and simulate BGP operations according to RFC4271 [3]. More precisely, we maintain a routing table for each peer – similar to BGP's Adj-RIB-In – the *adjacent routing table*. Each route announced by a peer is added to that peer's adjacent routing table. Whenever a withdrawal is received for a prefix, every route to that prefix is removed from the peer's adjacent routing table. Since we are not interested in routing traffic, we do not try to select preferred routes. We are, however, interested in knowing if a prefix p is up, i.e. if p is present in any of the adjacent routing tables.

The set of origins $\mathcal{O}(p)$ associated with prefix p is composed of the union of all the origins included in all of the AS paths of each adjacent routing table. If the cardinality of $\mathcal{O}(p)$ is larger than 1, p is a MOAS. For example, in Fig. 1, during $]t_0, t_1[$, $\mathcal{O}_{]t_0,t_1[} = \{1\}$ whose cardinality is 1, and the prefix is a SOAS. During $]t_1, t_2[$, $\mathcal{O}_{]t_1,t_2[}(p) = \{1, 2\}$ whose cardinality is 2, and the prefix is a MOAS. Finally, during $]t_3, t_4[$, $\mathcal{O}_{]t_3,t_4[}(p) = \varnothing$ whose cardinality is 0, and the prefix is down.

2.3 Methodology

Our study starts by revisiting the previous works [2, 5]: we compare the uptimes of MOAS and SOAS prefixes, and put into perspective the average uptime of MOAS prefixes with the average duration of MOAS events. By doing this, we show that our results are similar to what had been observed by [2, 5], thus ensuring that our further results can be put into perspective with the results provided by both of these studies.

We then consider that MOASes are not only a set of independent events, but they are related to a prefix. By grouping MOAS events per prefix, we are able to uncover inner relationships between multiple successive events. Most notably, we show that a large fraction of short-lived MOAS events are not the result of misconfigurations, which contradicts [2].

We look at topology graphs of MOAS prefixes from which we extract MOAS patterns that we use to classify and quantify MOASes. We study the temporal evolution of the topology graphs related to MOAS prefixes by comparing them months before and after MOAS events. This evolution enables us to understand the root causes behind the MOAS patterns.

3 Previous Work

Zhao et al. [2] pioneered the analysis of MOAS events, and analysed BGP data between late 1997 and mid 2001. This analysis concluded that 36% of the MOAS events were one-time events and lasted less than a day, 30% of which were attributed to a single misconfiguration. Excluding those, the average MOAS duration was 30.9 days. For MOASes that lasted over 9 days, the mean duration was 107.5d. These figures are computed using the MOAS duration per event.

The authors then discuss a number of reasons for which a prefix would be originated from multiple ASes: prefixes associated with an Internet exchange point (IXP) may be advertised by all the ASes within the IXP, since they are reachable through all of them. Multihoming without BGP (i.e. via static links or some IGP protocol) also leads to MOAS, since the prefixes are then announced by the upstream providers. Multihoming with BGP, but with a private ASN yields the same result. Anycasting can also lead to MOAS prefixes. Finally, since prefix aggregation in BGP transforms the AS path into an AS set (in which the order of ASNs is random), some artificial MOAS prefixes can be observed.

Chin [5] revisited the work of Zhao et al. by studying three weeks of data in January 2007, and found an average lifetime of MOAS events to be 13.25 hours. Chin then proposed new reasons behind MOAS prefixes: multinational companies may advertise prefixes from various branches in different countries, and such organizations possibly own multiple AS numbers. Companies may also host their servers in data centers, announcing the prefixes both, from the data center and from their offices.

Of course, MOAS prefixes can be the result of a malicious attack against the routing infrastructure, in which case they are often referred to as MOAS conflicts. As a result, the problem has been widely discussed in the literature related to prefix hijacking, such as [6, 7]. However these hijack papers usually focus on the threat posed by MOASes, not on their characteristics or classification.

4 Results

4.1 General Results

During the year of 2002, almost 310k different prefixes were announced, less than 9% of which presented (at least) one MOAS event. In 2012, there were almost 765k distinct announced prefixes, less than 6% of which were in a MOAS state at some point during the year. These figures suggest that, while both, the number of global prefixes and the number of MOAS prefixes increased in 10 years, their proportion has decreased. In

		uptime			lifetime		
		μ	CoV	q_{50}	μ	CoV	q_{50}
MOAS	2002	328d	0.25	363d	334d	0.23	364d
	2012	308d	0.34	364d	317d	0.31	364d
SOAS	2002	146d	1.11	37d	172d	0.89	146d
	2012	223d	0.72	348d	239d	0.65	360d

Table 1: General statistics on BGP data for 2002 and 2012

		μ	CoV	q_{50}
All MOAS events	2002	33d	2.23	22h
	2012	48d	1.88	26h
Short-lived MOAS events	2002	133mn	2.26	9.3mn
	2012	101mn	2.60	3.13mn

Table 2: Duration of MOAS events

both cases, less than 5% of MOAS prefixes were the result of route aggregations. We removed these prefixes from our MOAS cases before further analysis.

Table 1 shows the mean (μ), coefficient of variation (CoV = stdev / μ), and median (q_{50}) durations for the uptime and the lifetime of both MOAS and SOAS prefixes during 2002 and 2012. Mean values for MOAS prefixes in both, 2002 and 2012 are significantly higher than the values for SOAS prefix in terms of uptime and lifetime. This suggests the use of MOAS to improve the connectivity of a prefix. In particular, median uptime and lifetime of MOAS prefixes are both close to 1 year, meaning that 50% of those prefixes were seen over the entire observation period. The mean and median value for SOAS prefixes in 2012 – both close to 1 year – are also much higher than those in 2002, where the median uptime of 37d is very low compared to the observation period of 1 year, and to a median lifetime of 146d. These figures for 2002 are in line with the ones presented in [8], even though their analysis does not focus on MOASes, which strengthens our confidence in the accuracy of our method. While we only detail 2002 and 2012, we looked at the data of years in between and found similar conclusions.

MOAS Events In this section, we consider MOAS events as a set of distinct events, independent of the prefix with which they are associated. For example, we consider independently the 3 MOAS events depicted in Fig. 1 during $]t_1, t_2[$, $]t_5, t_6[$, and $]t_7, t_8[$. The **MOAS duration** (per event) is the duration of a single event. In Fig. 1, the durations of the three MOASes are $t_2 - t_1$, $t_6 - t_5$, and $t_8 - t_7$.

Figure 2 depicts the duration of MOAS events in 2002 and in 2012. MOAS duration information for 2002 and 2012 are available in Table 2. The large difference between the mean and the median shows how prevalent short-duration events are. The consequence of the comparison between these values and those presented in Table 1 is that MOAS prefixes do not spend their whole life in a MOAS state.

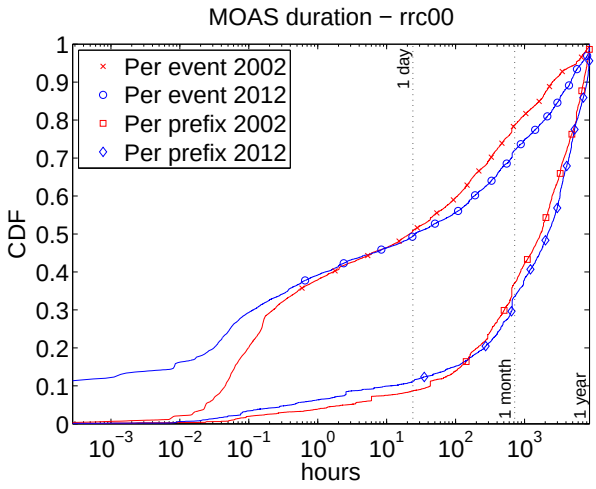

Fig. 2: MOAS events duration

MOAS Prefixes In this section, we consider MOAS events grouped by the prefix for which they appeared. Distinct MOAS events may appear over the course of the observation period for a single prefix p. We say two MOAS events associated with a prefix p are **distinct** if the origin sets $\mathcal{O}(p)$ are different for the two events. For example, in Fig. 1, prefix p has 3 MOAS events: during $]t_1, t_2[$, $]t_5, t_6[$, and $]t_7, t_8[$. Moreover, $\mathcal{O}_{]t_1,t_2[}(p) = \mathcal{O}_{]t_7,t_8[}(p) \neq \mathcal{O}_{]t_5,t_6[}(p)$. So, even though Fig. 1 depicts 3 MOAS events, only 2 of them are distinct in the sense that they involve different ASes. Furthermore, the **duration** of **MOAS events per prefix** is the sum of the durations of the individual MOASes associated with this prefix. Using Fig. 1, the MOAS duration for prefix p is $t_2 - t_1 + t_6 - t_5 + t_8 - t_7$. In the remainder of this section, unless explicitly stated, duration means the duration of the MOAS events *per prefix*.

Figure 2 plots the duration of MOAS events per prefix. Only around 10% of the MOASes are short-lived, which heavily contrasts with the 50% obtained when considering each MOAS event on its own. This implies that certain prefixes must have many MOAS events. This is confirmed by Fig. 3, where the number of MOAS events and the number of distinct MOAS events per prefix are plotted[3]. The prefixes are sorted by decreasing number of MOAS events. For the first 1000 prefixes with the most MOAS events, the mean and median duration of single MOAS events is very small (in the order of a few minutes or less).

Figure 3 shows that, for approximately 1000 prefixes out of the 43k MOAS prefixes, the number of *distinct* MOASes is significantly lower than the number of MOAS events. Some of these prefixes only have 1 distinct MOAS, but hundreds of MOAS events. In these cases, there was a continuous flipping between SOAS and MOAS announcements. This kind of behaviour can be explained by an instability between the prefix owner and

[3] The equivalent figure for 2002 looks very much alike, and was not included in the paper due to space restrictions.

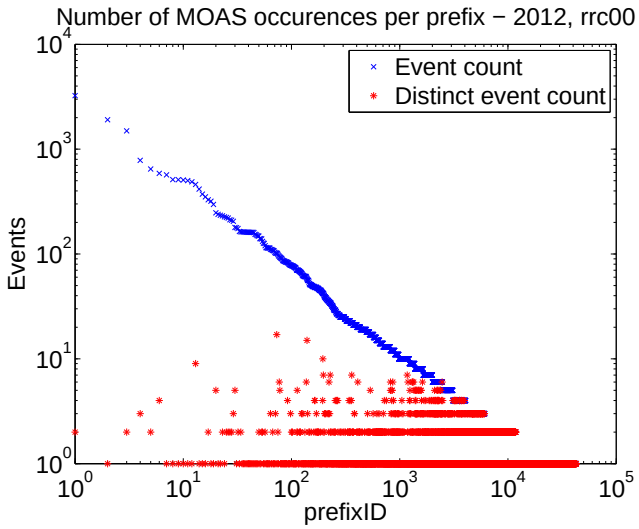

Fig. 3: Number of MOAS events per prefix

one of its upstreams. However, by looking at the AS paths in the duplicate BGP update messages related to these events, we saw that only one sub-path actually caused this flipping phenomenon to the route collector. For this reason, we suspect that this flipping was not caused by an instability in the prefix owner's connections, but by some router located in an AS between the collector and the prefix origin.

Figure 2 also implies that the bulk of short-lived MOAS events *cannot* be attributed to misconfigurations. If, as supposed by [2, 5], most short-lived MOAS events are due to a misconfiguration, there should not be that many recurrent MOAS events for the same prefix. Indeed, only the sum of numerous short-lived events for the same prefixes (Fig. 3) can result in raising the MOAS duration per prefix as much, compared to the MOAS duration per event. Since misconfigurations are usually sorted out promptly [9], many short-lived events affecting many distinct prefixes would not shift the CDF plot to the right. As a result, the curves in Fig. 2 would not show such a drastic difference between the "per event" and the "per prefix" computations.

The mean and median MOAS durations per prefix in 2002 and 2012 are detailed in Table 3. We clearly see that both the mean and median values for the MOASes per prefix are a lot larger than individual MOAS event durations. This is, once again, the result of the combination of the many short-lived events per prefix.

We also considered the fraction of MOAS uptime for a prefix over its total uptime. One might expect MOAS prefixes to remain in MOAS state during most of their uptime in order to maximize the benefits behind their chosen MOAS configuration. However, the distribution of the fraction of time in MOAS state distribution is uniform and contradicts this expectation. We explain this phenomenon by the use of *transient* MOAS configurations. A temporal analysis of the topological evolution of MOAS networks uncovered multiple cases of stub networks switching between upstream AS providers.

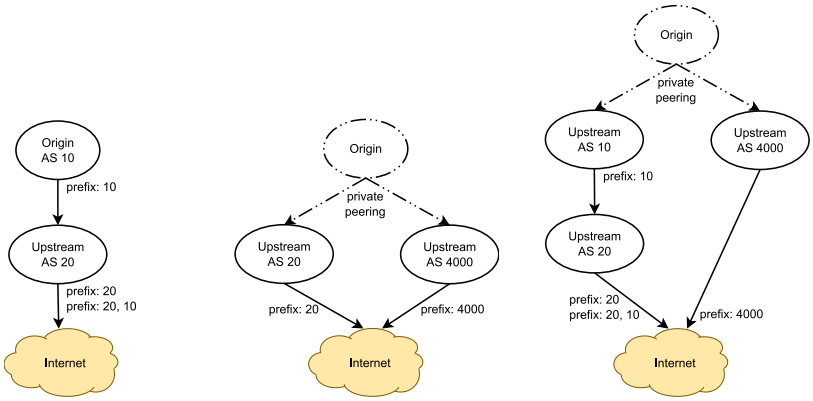

(a) Peering MOAS (b) Classical MOAS (c) Me-Too MOAS

Fig. 4: Graphs of MOAS patterns

This operation can be summarized as follows. Originally, prefix p is announced by ISP A. At some point, the owners of p find it more advantageous to use ISP B. In order to avoid any service disruption, p remains connected to (and announced by) A while things are being set up with B (i.e. connecting p to B), and then also starting announcing it from B. This results in a MOAS. After some time (weeks), p is disconnected from A, and remains exclusively announced and reachable via B.

4.2 MOAS Patterns

We analyzed the AS-level graph of MOAS prefixes and were able to extract a set of patterns that result from MOAS announcements. This led us to a taxonomy of MOASes that we present now. In order to understand the reasons behind these MOAS events, we looked at the evolution of the AS topology of MOAS networks during 6 months surrounding the MOAS occurence.

The first pattern, depicted in Fig. 4a, shows a situation where both, the prefix owner and its upstream are announcing the prefix. We call this situation a **Peering MOAS**. Even though Fig. 4a only depicts one upstream, we saw cases where upstreams of the upstream were also announcing that prefix. The mean and median durations for peering MOASes are presented in Table 3. Figure 5 shows the distribution of the durations of peering MOASes for 2002 and 2012. 60% of those events last longer than a month, and around 10% of them are short-lived.

We often saw this pattern appearing in the following setting. The prefix is first announced by the upstream, but assigned to the customer (e.g. Fig. 1, during $]t_0, t_1[$). At some point, the customer decides to handle routing on its own, and acquires its own AS number and starts BGP peering with the upstream. At this point, there is a MOAS (Fig. 1, during $]t_1, t_2[$). Eventually, the upstream withdraws its announcement of the prefix, leaving only the owner's announcement in the routing tables (Fig. 1, during

		μ	CoV	q_{50}
All MOAS prefixes	**2002**	111d	1.01	71d
	2012	125d	0.95	90d
Peering MOAS pattern	**2002**	103d	1.03	64d
	2012	123d	0.97	88d
Classical MOAS pattern	**2002**	80d	1.24	32d
	2012	101d	1.12	43d
Me-Too MOAS pattern	**2002**	203d	0.64	241d
	2012	181d	0.60	197d

Table 3: Duration of MOAS prefixes

$]t_2, t_3[$). In this case, the MOAS was the side-effect of a real topology change. Even though we did not explicitly witness any situation in which this description is not accurate, this pattern is not necessarily the result of a provider-customer relationship. It is conceivable for this pattern to be the result of any direct peering relationship, such as peer-to-peer or siblings networks.

In other peering MOAS cases, the owner has several upstream providers, some of which re-announce the prefix. Effectively, the prefix owner is multihomed; but a subset of its upstreams originate the prefix. We believe that in this situation, the network was originally connected to a single upstream that handled BGP operations on its behalf, but then decided to multihome in order to benefit from increased connectivity. However, the original ISP's configuration remains unchanged and carries on announcing the prefix.

We consider this pattern to create **fake** MOASes because, in both of these situations, there is no gain for the owner from its upstream's announcement. Indeed, if the upstream stopped originating the prefix, the situation would remain unchanged: the prefix would still be reachable via all of its upstreams without loss of connectivity. Table 4 shows that peering MOASes amount to around 70% of all MOAS events. We believe this class of MOAS is caused by loosely-defined (or outdated) routing policy. Another cause would be prefixes associated with IXPs, as described by [2, 5].

The second pattern, depicted in Fig. 4b, is the expected AS pattern when talking about MOASes. For this reason, we call it the **classical MOAS** pattern. There are multiple distinct AS paths leading to the prefix. The mean duration of these MOASes are shown in the penultimate row of Table 3. These values suggest that classical MOASes are longer-lived in 2012 than in 2002. This is confirmed by Fig. 5 which plots the durations of these events. In 2002, around 50% of them were short-lived, which then decreased to around 35% for 2012.

We found the main reasons behind this pattern to be in accordance with engineering practices described in [2, 5]. In order to verify this, we used WHOIS data for the prefix and origin ASes. The ASes most often belonged to well-established ISPs, and the prefixes were registered to another entity. We also saw cases where multinational companies were the owner of each of the origin ASes.

Table 4 shows the proportion of classical MOASes among all MOASes, which is around 25%. Consequently, the pattern that is traditionally believed to be *the* MOAS configuration only amounts to a quarter of MOAS prefixes.

Any loss of origin in a classical MOAS means a loss of connectivity between the prefix and its upstream. If an origin AS stops announcing the prefix, it will not receive traffic for it. It will therefore not provide any connectivity to the Internet for the owner. This contrasts with the situation of multihoming with a fake/peering MOAS, where the loss of an upstream origin does not affect the connectivity of the network, since the upstream AS remains in the AS path to the origin.

The last pattern, depicted in Fig. 4c, is named **Me-Too MOAS** to underline its "being over-announced" property. It is composed of both of the previous patterns at a single time: the left-hand side of Fig. 4c shows a peering MOAS, while the first-level AS peers are arranged in a classical MOAS manner. The mean and median durations of this pattern is shown in the last row of Table 3. These values suggest that me-too MOASes are stable. Figure 5 confirms that few of these events are short-lived (less than 5% in both cases), and over 80% of them last longer than two months.

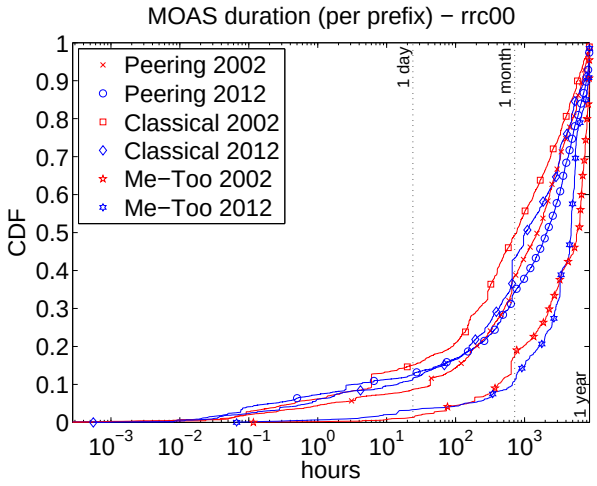

Fig. 5: MOAS patterns duration (per prefix)

We saw this pattern appear in the two following situations. The first one was a combination of subletting of IP space. Using Fig. 4c as illustration, the prefix block p is owned by AS20 and AS10 rents it. The WHOIS record associated with p clearly stated that prefix p was part of non-transferable IP addresses. So, because AS20 is the owner, it keeps on announcing p. However, since AS10 rents it, it also announces the prefix. This results in a peering MOAS, i.e. the left-hand side of Fig. 4c. Additionally, AS10 assigned p to one of their customer for use. At some point, this customer chooses to do multihoming and uses AS4000 for that purpose. In return, AS4000 announces p as

well, i.e. the right-hand side of Fig. 4c. The second situation was when a prefix owner decided to change upstreams. Originally, the owner's prefix was announced by a tier-1 ISP which used multiple AS numbers, one for its global activities (AS20 in Fig. 4c), and one for its local activities (AS10 in Fig. 4c). However the ISP used both of those AS numbers to originate the prefix, although it needs to go through the local AS from the backbone to reach the customer (this corresponds to a peering MOAS). Then, the user (AS at the top of Fig. 4c) decides to switch their ISP service to another tier-3 ISP (AS4000 in Fig. 4c). During the transition, which usually lasts several weeks, the prefix was announced by both the old tier-1 (AS10 and AS20 in Fig. 4c) and the new local tier-3 ISP (Fig. 4c, AS4000). This situation, then, presents a peering MOAS with a classical MOAS.

Per prefix	**2002**	**2012**
Peering MOAS (a)	72.55%	72.63%
Classical MOAS (b)	31.09%	28.6%
Me-Too MOAS (c)	5.5%	3.84%
(a) & (b)	6.37%	3.24%
(a) & (c)	1.95%	1.59%
(b) & (c)	1.37%	0.49%
(a) & (b) & (c)	0.52%	0.24%

Table 4: Proportion of occurrences of MOAS patterns

Table 4 shows that me-too MOAS events amount to 3% to 5% of MOASes. This can be explained by the fact that this configuration is unlikely to arise from erroneous situations, unlike the previous two patterns since it requires (at least) 3 origin ASes for a single prefix, *with* a peering relation among two of them.

The bottom rows of Table 4 show the proportion of prefixes that exhibit different types of MOAS patterns. These values suggest that MOAS prefixes only exhibit one kind of MOAS event throughout their lifetime. When we put this information in relation with the MOAS durations in Table 3 (125d on average) and the MOAS prefix uptime in Table 1 (308d on average), it is clear that MOAS prefixes do not spend their whole uptime in a MOAS state. However, the fact that the MOAS prefixes do not switch from one MOAS class to another suggests that their configuration remains stable. We can think of two main reasons why these MOAS announcements would be withdrawn. A reason could be that the owner of the prefix intentionally withdraws this announcement, for exemple due to exceeding the bandwidth allowance of one of its peer. Another reason is the data bias from our collector router, i.e. these routes are not propagated to the collector anymore because they have been filtered out.

5 Conclusion

In this paper, we studied MOAS events in multiple ways. First, we revisited previous works by looking at MOAS events on their own. Then we considered MOAS events

along with their prefix. Grouping the events that way underlines the relationship between seemingly independent MOAS events. Most notably, we showed how many short-lived events repeat in order to result in a long-lived MOAS prefix. This observation eliminates the possibility that these events are the result of a misconfiguration.

We also looked at the evolution of the topology graph of MOAS prefixes and we classified MOASes into three distinct patterns. The most popular pattern, peering MOASes, is composed of long-lived MOASes where the different origin ASes are directly peering. We consider these as fake MOASes because there is no benefit from the MOAS announcement. The second class of MOAS is composed of classical MOASes. This is the standard MOAS configuration. However, they only make up for a third of the global MOAS events and global MOAS prefixes. The last pattern, me-too pattern, is a combination of the other patterns and was encountered as a transitional configuration when an owner was switching its upstream provider to another one.

Finally, we looked at data ten years apart, and showed that there is little difference in MOAS properties in terms of prefix uptime, MOAS duration, and MOAS classification/proportion. This is remarkable because the size of the network grew by around 400% over this period of time [10].

Future work includes expanding our ground-truth sources with verified peering information to supplement WHOIS data. This would permit further validation and classification of peering MOASes. Another direction is to deeper study the flipping between SOAS and MOAS related to a single prefix. As we suggested, it may be the result of an intervention of the owner, in order to comply to the terms of a peering agreement (e.g. exceeded bandwidth). Finally, our analysis makes use of a single vantage point to analyse MOAS conflicts. This certainly results in under-estimating the number of MOAS events seen, particularly in terms of peering MOASes. (Classical MOASes try, by design, to diversify the AS paths as much as possible.) Although we are confident that the global trends and orders of magnitudes we exposed in this study remain true regardless of the vantage point, using (multiple) different route collectors would certainly provide better estimates.

References

1. J. Hawkinson et al.: Guidelines for creation, selection, and registration of an Autonomous System (AS). RFC 1930 (March 1996)
2. X. Zhao et al.: An analysis of BGP multiple origin AS (MOAS) conflicts. In: IMW. (2001)
3. Y. Rekhter et al.: A Border Gateway Protocol 4 (BGP-4). RFC 4271
4. RIPE NCC: Routing Information Service. http://www.ripe.net/ris/
5. C. Kwan-Wu: On the characteristics of BGP multiple origin AS conflicts. In: ATNAC. (2007)
6. M. Lad et al.: PHAS: A Prefix Hijack Alert System. In: USENIX Security Symposium. (2006)
7. X. Shi et al.: Detecting prefix hijackings in the internet with argus. In: IMC. (2012)
8. G. Siganos et al.: BGP routing: A study at large time scale. In: Proc. IEEE Global Internet. (2002)
9. R. Mahajan et al.: Understanding bgp misconfiguration. SIGCOMM Comput. Commun. Rev. **32**(4) (August 2002) 3–16
10. G. Huston: BGP reports. http://bgp.potaroo.net/