

# Live Long and Prosper: Analyzing Long-Lived MOAS Prefixes in BGP

Khwaja Zubair Sediqi

Max Planck Institute for Informatics

zsediqi@mpi-inf.mpg.de

Anja Feldmann

Max Planck Institute for Informatics

anja@mpi-inf.mpg.de

Oliver Gasser

Max Planck Institute for Informatics

oliver.gasser@mpi-inf.mpg.de

**Abstract**—BGP exchanges reachability information in the form of prefixes, which are usually originated by a single Autonomous System (AS). If multiple ASes originate the same prefix, this is referred to as a Multiple Origin ASes (MOAS) prefix. One reason for MOAS prefixes are BGP prefix hijacks, which are mostly short-lived and have been studied extensively in the past years. In contrast to short-lived MOAS, long-lived MOAS have remained largely understudied.

In this paper, we focus on long-lived MOAS prefixes and perform an in-depth study over six years. We identify around 24k long-lived MOAS prefixes in IPv4 and 1.4k in IPv6 being announced in January 2023. By analyzing the RPKI status we find that more than 40% of MOAS prefixes have all origins registered correctly, with only a minority of MOAS having invalid origins. Moreover, we find that the most prominent CIDR size of MOAS prefixes is /24 for IPv4 and /48 for IPv6, suggesting their use for fine-grained traffic steering. We attribute a considerable number of MOAS prefixes to mergers and acquisitions of companies. Additionally, more than 90% of MOAS prefixes are originated by two origin ASes, with the majority of detected origin AS relations being customer-provider. Finally, we identify that the majority of MOAS users are IT companies, and just 0.9% of IPv4 MOAS and 6.3% of IPv6 MOAS prefixes are used for anycast.

## I. INTRODUCTION

To exchange reachability information about IP addresses, different Autonomous Systems (ASes) in the Internet use BGP [1]. In BGP, each IP address prefix (i.e., collection of IP addresses) is usually originated by a single AS [2]. There are, however, also cases where multiple ASes originate the *same prefix*, this prefix is in turn called a Multi Origin AS (MOAS) prefix. Network operators use MOAS prefixes to e.g., provide resilience, load balancing, and multi-homing.

In addition to these uses, MOAS prefixes can also be a result of mergers of companies operating two ASes, misconfigurations [3], and—most problematically—prefix hijacks. This last case occurs if an attacker hijacks traffic to a specific prefix by announcing this prefix with its own AS as an origin. Unfortunately, prefix hijacks are happening relatively frequently [4], [5]. Consequently, prefix hijacks have been the focus of numerous studies over the past years [6]–[13].

Surprisingly, the use of long-lived MOAS prefixes—i.e., when they are visible for a longer time in BGP—remains understudied.

When analyzing RIB snapshots of all RIPE RIS [14] and Routeviews [15] route collectors, we find that the number of




Fig. 1. Number of long-lived MOAS prefixes and origin ASes over time.

long-lived MOAS prefixes (i.e., visible for at least 30 days) is growing over time, as can be seen in Figure 1, both in terms of prefixes and origin ASes.

IPv4 long-lived MOAS prefixes increase from 10k in 2017 to over 24k prefixes at the beginning of 2023, with the number of origin ASes growing by about 50% in the same time period. Moreover, the fraction of long-lived MOAS prefixes out of all visible BGP prefixes is increasing as well (see Section III-D). In this paper, we focus on the analysis of these long-lived MOAS prefixes. We perform a longitudinal analysis across multiple dimensions to unveil prefix and origin characteristics, visibility, and users of MOAS prefixes. More specifically, the main contributions of this paper are:

- **Methodology to Detect Long-Lived MOAS Prefixes:** We apply a rigorous methodology to detect long-lived MOAS prefixes over a period of six years (see Section III). We use the Kneedle algorithm to discern long-lived from short-lived MOAS prefixes and perform a sensitivity analysis to analyze the influence of route collector artifacts.
- **MOAS Prefixes and Origins:** We perform a longitudinal analysis of MOAS prefixes and origin ASes (see Section IV). We find that more than 40% of long-lived MOAS prefixes have a valid RPKI status for all origins, hinting that hijacks are not prevalent among them. Moreover, we identify that the vast majority of MOAS prefixes are announced by two origin ASes. The most commonly used CIDR sizes are /24 for IPv4 and /48

for IPv6, suggesting their use for fine-granular traffic steering. Additionally, we find a considerable number of MOAS prefixes linked to mergers of companies.

- **Visibility Analysis:** We analyze the visibility of MOAS prefixes across route collectors (see Section V). Most MOAS prefixes are visible on hundreds of route collector (RC) peers with one origin AS, whereas the other origin AS is visible only on a handful of RC peers.
- **Users and Usage of MOAS Prefixes:** To better understand users and usage of MOAS prefixes we do an in-depth analysis on the AS relationship, their business types, and the prevalence of anycast within MOAS prefixes (see Section VI). Our results show that the majority of origin ASes are used by IT companies being in a customer-provider relationship, and only a small fraction of MOAS prefixes are used for anycast.

## II. DATASETS

To analyze the prevalence and evolution of MOAS prefixes in the Internet, we use the following datasets.

**BGP Datasets:** To get a good understanding of MOAS prefixes in the Internet, we use route collectors (RCs) from two different route collector projects: RIPE RIS [14] and Routeviews [15]. We use 59 route collectors throughout our six years study period from January 1, 2017, to January 1, 2023. During the study period, we analyze daily routing information base (RIB) snapshots of all RCs for the duration of 2098 consecutive days. To calibrate our methodology (see Section III) we compile an auxiliary BGP dataset by using three RIBs per day for a period of six months in 2022 from the same RCs.

**AS Relationships Dataset:** We use CAIDA’s AS relationship dataset [16] from January 1, 2017, to January 1, 2023, to analyze inter-AS relationships of origin ASes associated with MOAS prefixes. This dataset is released by CAIDA every three months. The dataset classifies the inter-AS relationships into “customer-provider or provider-customer”, “peer to peer”, or “unknown” using the method described by Luckie et al. [17].

**AS to Organizations Mapping Dataset:** We use the inferred AS to organization mapping dataset provided by CAIDA [18] for the duration of our study from January 1, 2017, to January 1, 2023. The dataset is created by querying WHOIS databases from the five Regional Internet Registries (ARIN, LACNIC, RIPE, AFRINIC, and APNIC), and two National Internet Registries (KRNIC and JPNIC) every quarter, and then applying CAIDA’s AS to organization technique to map each AS to an organization [19]. Due to the lack of historical data, we are unable to use two recently presented improvements [20], [21] to the CAIDA AS to organization dataset.

**ASdb Dataset:** To gather detailed information about the type of organization using MOAS prefixes, we use the ASdb dataset [22]. ASdb provides a mapping of each AS into 17 industry categories and 95 sub-categories using data from established business intelligence databases and machine learning. We use ASdb to identify AS organization types and business relationships between the origin ASes of MOAS prefixes.

**RPKI Dataset:** We use archived RPKI records from RIPE NCC [23] in a monthly granularity from January 2017 to January 2023. With the RPKI dataset we identify the route origin validation (ROV) state of MOAS prefixes.

**Anycast IP Prefix Dataset:** Moreover, we use a list of IPv4 and IPv6 prefixes classified as anycast prefixes by the BGP.Tools project [24]. The project uses an approach similar to MAnycast<sup>2</sup> [25] to run periodical anycast detection measurements. We use the published list of anycast IPv4 and IPv6 prefixes for the most recent date for our study, January 1, 2023, to analyze the prevalence of anycast deployments within detected MOAS prefixes.

## III. METHODOLOGY

In this study, we focus on long-lived MOAS prefixes, i.e., prefixes that are announced from multiple origin ASes for a substantial period of time. In contrast to prefix hijacks [6]–[13], such MOAS prefixes have received little attention from the research community in the past.

To identify long-lived MOAS prefixes and assess the impact of RC artifacts (e.g., missing RIBs, offline RC peers), we apply the following methodology: First, we extract all prefixes from the BGP dataset, apply a series of filters, and select prefixes appearing as MOAS. Next, we measure the lifetime of MOAS prefixes and perform a sensitivity analysis. Afterwards, we examine the effect of using more than one RIB snapshot per day in our pipeline. Finally, we identify and apply a lifetime threshold to distinguish between short-lived and long-lived MOAS prefixes. All these steps are detailed in the following.

### A. Extracting, Filtering, and Selecting MOAS Prefixes

To develop our methodology, we download and use daily RIB snapshots from all 59 available RCs of the RIPE-RIS [14] and Routeviews [15] projects for a six-months period spanning from June 1, 2022, to November 30, 2022. Since we focus on long-lived MOAS prefixes, we do not use BGP UPDATE files, as those provide only additional short-term information. As a first step, we extract all prefixes from RIB snapshots of RCs and apply a series of filters to remove the following artifacts: private or reserved origin AS numbers [26], reserved IP prefixes, prefixes that should not be globally routed, and special purpose IPv4 and IPv6 prefixes [27]–[29]. Furthermore, we remove default route (i.e., 0.0.0.0/0) or prefixes with network bits set beyond their CIDR size (i.e., 1.2.3.0/16). As a result of applying filters, we remove around 1% of all prefixes from the BGP dataset.

Then we check the origin AS of every IP prefix across all RIB snapshots at the exact date and time<sup>1</sup>, select prefixes with more than one origin AS, and consequently mark them as MOAS prefixes for further analysis.

<sup>1</sup>We download available RIBs daily at 08:00:00 UTC from all route collectors. It is important to check the RIB snapshots for the exact date and time across all route collectors, otherwise a prefix could be falsely classified as MOAS appearing with a different origin ASes at two different points in time.

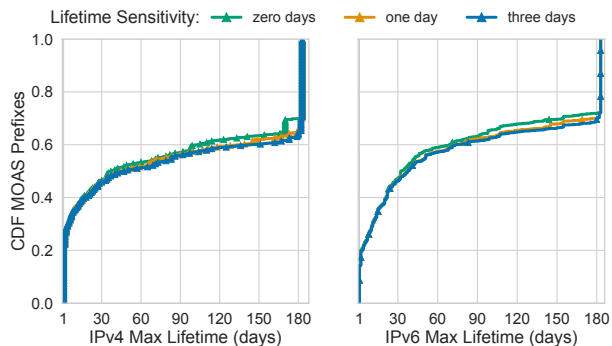


Fig. 2. CDF of maximum lifetime with various sensitivity thresholds for IPv4 (left) and IPv6 (right) MOAS prefixes.

### B. MOAS Lifetime Sensitivity Threshold

As BGP data from RC projects can be incomplete, e.g., due missing RIBs or RCs temporarily being offline, calculating the lifetime of a MOAS prefixes requires us to be extra cautious. We use the lifetime, i.e., the duration a prefix is seen as a MOAS continuously in RIB snapshots, as a metric to identify and select long-lived MOAS prefixes. For each MOAS prefix observed in the BGP dataset, we calculate the maximum lifetime over the six months period. To assess the effect of BGP dataset incompleteness, we perform a sensitivity analysis of the maximum lifetime using different sensitivity thresholds. We investigate how the sensitivity threshold—i.e., the maximum number of continuous days a MOAS prefix can be missing from RIB data before its lifetime ends—affects the maximum lifetime of MOAS prefixes.

In Figure 2, we show the maximum lifetime for IPv4 and IPv6 MOAS prefixes for sensitivity values of zero, one, and three days. We find that more than 40% of IPv4 and around 50% of IPv6 MOAS prefixes have a maximum lifetime of fewer than 30 days. On the other end, around 35% of IPv4 and 30% of IPv6 MOAS prefixes are visible throughout the entire six months. Moreover, around 25% of IPv4 and 20% of IPv6 MOAS prefixes have a maximum lifetime of at most five days.

Looking at the effect of different sensitivity values, we notice that overall the one and three days sensitivity values slightly increase the lifetime of MOAS prefixes compared to the zero days sensitivity. This difference is most visible for IPv4 in the interval of 150 to 180 days, where we see the effect of a route collector outage when the RIB snapshot of Sydney RC for time 08:00:00 UTC is missing on June 13, 2022, but available for the rest of the time. Considering this and as a precaution to mitigate possible effects of unavailable RCs or RC peers on the lifetime of MOAS prefixes, we use the one day sensitivity threshold in the remainder of the paper. We do not consider the three days sensitivity threshold, because it does not substantially increase the lifetime of MOAS prefixes beyond the one day threshold, and a gap of three days during which a MOAS prefix is not visible is large enough to possibly




Fig. 3. CDF of maximum lifetime (left) and observability (right) of MOAS prefixes for single and three RIBs per day.

lead to incorrect calculations of MOAS prefix lifetimes.

### C. Impact of Using More Than One RIB per Day

After having performed the lifetime analysis with various sensitivity values, we examine if increasing the number of RIB snapshots per day increases the lifetime of MOAS prefixes. Therefore, we use three RIB snapshots per day (3RIBs/24hours), as this is the maximum number of possible RIB snapshots with the same timestamp across both Routeviews and RIPE RIS RCs. We study the effect of increasing the number of RIB snapshots considering the following two scenarios.

First, we examine by how much the increased number of RIBs increases the maximum lifetime of MOAS prefixes. The left subplot of Figure 3 illustrates the comparison and effect of using three RIBs compared to one RIB per 24 hours on the maximum lifetime of MOAS prefixes. As can be seen, increasing the number of RIBs to three per day (blue line) does not substantially increase the maximum lifetime of MOAS prefixes compared to a single RIB per day, orange line. Furthermore, the number of MOAS prefixes visible throughout the complete six months remains similar. The higher share of MOAS prefixes with a maximum lifetime of fewer than five days in three RIBs compared to a single one is because we observe many more short-lived MOAS (e.g., MOAS appearing only once) when using three RIBs per day.

Second, we check if increasing the number of RIB snapshots increases the observability of MOAS prefixes. We define the observability of a MOAS prefix as the number of days out of the total days, when a prefix is observed as a MOAS. Therefore, the observability metric allows us to see how consistently prefixes are visible as MOAS. As shown in the right subplot of Figure 2, for nearly 80% of MOAS both one and three RIBs have more than 95% observability. This high value gives us confidence that using the maximum lifetime of a MOAS prefix is a good approximation of its actual lifetime. Moreover, with three RIBs we see a higher fraction of MOAS with 100% observability compared to one RIB. This is because with three RIBs we observe more MOAS prefixes that are

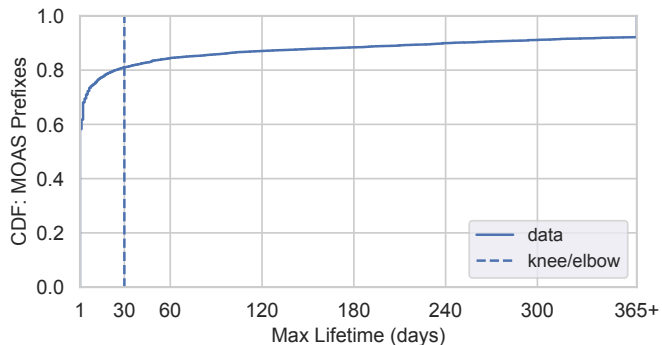


Fig. 4. CDF of maximum lifetime of IPv4 and IPv6 MOAS prefixes capped at 365 days, vertical line shows the elbow/knee of the distribution.

visible only on a single RIB snapshot, which by definition have an observability of 100%.

To summarize, increasing the number of RIBs from one to three per day does not increase the lifetime or observability of MOAS prefixes, particularly long-lived ones. Instead, three RIBs per day add more short-lived MOAS, which are not the focus of this study. Therefore, in the remainder of the paper we use one RIB per day.

#### D. Identifying Long-Lived MOAS Prefixes

Previous studies have shown that prefix hijacks, route leaks, and configuration errors cause short-lived MOAS [30]–[34]. We strive to analyze long-lived MOAS prefixes in BGP. Therefore, we develop a technique to identify a threshold to distinguish between short-lived and long-lived MOAS prefixes.

We measure the maximum lifetime of all MOAS prefixes for our complete BGP dataset as observed from January 1, 2017, to January 1, 2023. Then, we use the mathematical Kneedle algorithm [35] to determine the “elbow” of the maximum lifetime of MOAS prefixes. The “elbow” point identifies the maximum curvature value within the lifetime of all MOAS prefixes and we use it as distinction point between short-lived and long-lived MOAS prefixes.

Figure 4 shows the maximum lifetime (capped at 365 days) for all MOAS prefixes. Using the Kneedle algorithm we identify 30 days as the elbow, which we therefore use as a threshold to distinguish between short-lived and long-lived MOAS prefixes. Generally, we find that more than 80% of MOAS prefixes are short-lived, whereas less than 20% have a lifetime of at least 30 days.

Now that we defined the threshold of minimum 30 days lifetime between short-lived and long-lived MOAS prefixes, we examine the overall share of MOAS prefixes in the BGP dataset. In Figure 5 we show the fraction of MOAS prefixes among all prefixes for the duration of six years for IPv4 and IPv6, respectively. The fluctuation in the number of IPv4 MOAS prefixes for the upper subplot is due to Angola Cables (AS37468) creating MOAS conflicts for around 90k prefixes, on May 12, 2018, leading to the 10% IPv4 MOAS peak.

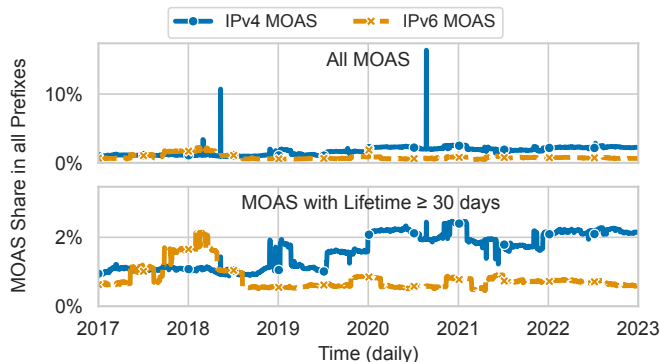


Fig. 5. Fraction of MOAS prefixes out of all prefixes for IPv4 and IPv6 over time, all MOAS (top) and long-lived MOAS (bottom).

Moreover, on August 24, 2020, Huge Networks - DDoS Mitigation (AS264409) creates MOAS conflicts for 143k prefixes, resulting in the 15% peak of IPv4 MOAS prefixes.

The lower subplot is limited to long-lived MOAS prefixes only, and exhibits less variation. We see the fraction of MOAS prefixes in IPv4 increasing from around 1% in 2017 to around 2% in 2023. In IPv6, the fraction is relatively stable at just below 1%, with a peak in February 2018. This peak is a result of GTT Communications (AS3257, AS5580) and DigitalOcean (AS14061, AS200130, AS202018, AS202109, AS393406, AS62567) each announcing hundreds of IPv6 hyper-specific [36] (i.e., /49–/128) MOAS prefixes.

In the remainder of the paper we focus exclusively on long-lived MOAS prefixes. Therefore, we now use the term MOAS prefix interchangeably with long-lived MOAS prefix.

#### E. Limitations

While we try to address multiple caveats with our approach, it still has some limitations. Our observation of MOAS prefixes in the BGP dataset is limited to peer ASes of RCs. Therefore, our detected MOAS prefixes should be seen as a *lower-bound*, i.e., there might be more MOAS prefixes out there which are not visible at route collectors. As the vast majority—95% for IPv4 and 91% for IPv6—of MOAS prefixes are originated by the exact same set of origins throughout the six months, we do not give special treatment to MOAS prefixes with changing AS origins over time. Despite these limitations, we argue that our approach is suitable and our analysis provides valuable insights into the prevalence, characteristics, and use of MOAS prefixes in the Internet.

## IV. PREFIXES AND ORIGINS

In this section we study MOAS prefixes and origins in detail: We analyze the RPKI status, CIDR sizes, and origin ASes of MOAS prefixes.

#### A. RPKI Status of MOAS Prefixes

As MOAS prefixes by definition are announced by at least two origin ASes, we now investigate whether a MOAS prefix has valid RPKI entries for its origin ASes. We examine the

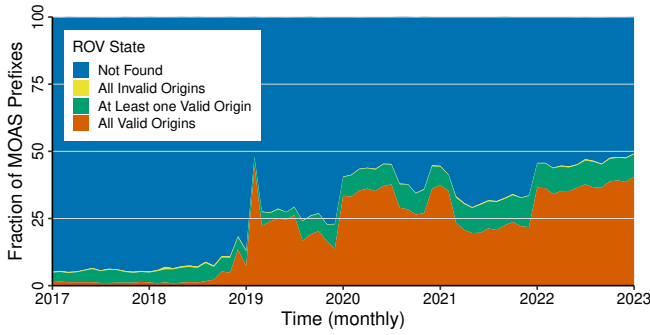


Fig. 6. MOAS Route Origin Validation state in the RPKI dataset over time.

route origin validation (ROV) state of MOAS prefixes using the RPKI dataset [23] for all origins—i.e., for every prefix-origin (PO) pair—on a monthly basis, starting from January 2017 to January 2023. We classify the ROV state of MOAS prefixes as follows: (i) “All Valid Origins” if all PO pairs have a valid route origin authorization (ROA) record; (ii) “At Least One Valid Origin” if at least one PO pair of a MOAS prefix has a valid ROA record; (iii) “All Invalid Origins” if all PO pairs conflict with a present ROA record; (iv) “Not Found” if we do not observe any ROA for a MOAS prefix.

In Figure 6 we show the RPKI classification of MOAS prefixes over time. We note that the percentage of MOAS prefixes with “All Valid Origins” (shown in orange), has increased from less than 5% in 2017 to around 40% in January 2023. The increase in share of valid ROV states is in line with the RPKI deployment in all BGP prefixes [37]. It also underlines that a large part of MOAS prefixes are not due to prefix hijacks, route leaks, or router configuration errors, but are in fact intended for use. We see a peak in the beginning of 2019 for “All Valid Origins”, which reaches around 45% and is due to TTNNet (AS47331) and Turk Telekomunikasyon (AS9121) announcing nearly 8k IPv4 MOAS prefixes—mostly /24 and /21–/23—as a result of their merger [38]. The share of MOAS prefixes with at least one valid origin (shown in green) has slightly increased from 2017 to 2023. Manually investigating the partially valid cases, we notice that not all origin ASes entered the prefix information properly in the RPKI database. We find cases where both ASes are registered with same organization name, with one having a valid and the other having an invalid ROV state for the MOAS prefix. Another possible reason for partially valid MOAS prefixes are router misconfigurations. The chance of them being hijacked prefixes is relatively small, as hijacks usually have short durations [6]–[13] and our analyzed prefixes are seen for at least 30 days. MOAS prefixes with all invalid origins (shown in yellow), contribute less than 1% throughout our measurement period. The share of MOAS prefixes for which no ROA entry exists (shown in blue) decreases from 90% in 2017 to around 50% in 2023.

**To summarize:** Around 40% of MOAS prefixes in January 2023 have all their origins validated in RPKI, with another  $\approx 10\%$  having a partially valid RPKI status. The case of all

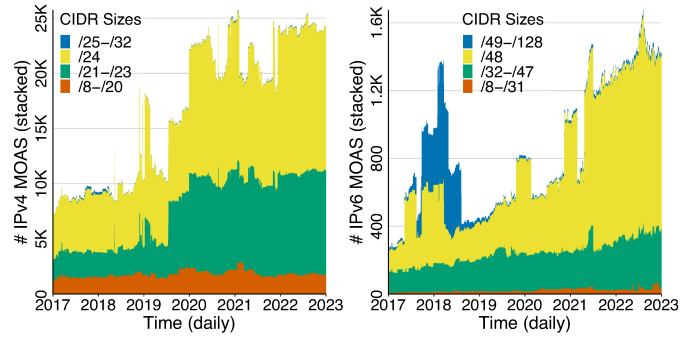


Fig. 7. MOAS prefix CIDR sizes over time, for IPv4 (left) and IPv6 (right).

origins being invalid occurs for less than 1% of cases. This underlines that a large part of long-lived MOAS prefixes is likely intended for actual use and not linked to hijacking, misconfigurations, or route leaks.

### B. CIDR Sizes

Prefixes in BGP are announced at different granularity levels, i.e., CIDR sizes. These CIDR sizes can provide insights into the usage of prefixes. Therefore, we analyze CIDR sizes of MOAS prefixes as shown in Figure 7 (IPv4 on the left, IPv6 on the right). For IPv4, we find that MOAS prefixes overall have increased from fewer than 10k in 2017 to more than 24k in 2023. The /24 CIDR size is most prominent in IPv4, followed by the /21–/23 CIDR size group. Hyper-specific [36] MOAS prefixes, i.e., /25–/32, constitute a tiny fraction of MOAS prefixes, as shown in blue.

While the number of MOAS prefixes with a CIDR size of /8–/20 is relatively stable over time, both the /24 and the /21–/23 range increase by more than 50% after mid-2019. We investigate the peaks in the data and attribute them to the merger of TTNNet and Turk Telekomunikasyon, see Section IV-A. These prefixes are visible as a substantial increase in the /21–/23 CIDR size group from August 2019 to January 2023. In addition, the peaks of /24 MOAS prefixes in July 2019, January 2020, September 2020, and January 2022 are associated with thousands of MOAS prefixes originating from Orange Spain (AS12479, AS12715). The latter of the two ASes used to belong to Jazztel, which was acquired by Orange in 2014 [39]. Moreover, the acquisition of KPN International (AS286) by GTT (AS3257) in December 2019 [40] still results in 389 /24 MOAS prefixes in January 2023.

In IPv6 we see an increase from fewer than 400 to more than 1400 MOAS prefixes, as shown in the right plot of Figure 7. The number of /48 MOAS shown in yellow, which forms the largest CIDR size in IPv6, increases more than five-fold from 2017 to 2023. The CIDR size ranges /32–/47 and /8–/31 form the second and third-largest group of IPv6 MOAS prefixes and grow steadily. We see a large number of IPv6 hyper-specific MOAS prefixes, i.e., /49–/128 CIDR sizes shown in blue, from mid-2017 to mid-2018, which are normally not routable in the Internet [36]. These are again linked to GTT and DigitalOcean, see Section III-D.




Fig. 8. Number of origin ASes per prefix (left) and the number of originated MOAS prefixes per origin AS set (right) for January 1, 2023.

**To summarize:** ASes use mostly fine-granular CIDR sizes—i.e., /24 for IPv4 and /48 for IPv6—to announce MOAS prefixes in the Internet, which allows for more fine-grained control of routing policies. Overall, the distribution of CIDR sizes within MOAS prefixes is similar to non-MOAS prefixes reported by BGP routing table [41]. Moreover, we find that a considerable number of MOAS prefixes are related to mergers and acquisitions of companies.

### C. Origin ASes

Next, we analyze the origin ASes for every MOAS prefix in more detail. In the left subplot of Figure 8 we observe that more than 95% of IPv4 and around 88% of IPv6 prefixes have only two origin ASes (i.e., the minimum required for them to be classified as a MOAS prefix). Moreover, around 1% of IPv4 and 8% of IPv6 MOAS prefixes have more than three origin ASes. We notice that a few prefixes are originated by more than 80 origin ASes. We manually investigate these cases and identify Verisign—a DNS provider which also operates two DNS root servers—to have the highest number of origin ASes for one IPv4 and one IPv6 prefix.

Furthermore, we analyze the number of MOAS prefixes originating from the same set of origin ASes, so-called origin AS sets. In Figure 8 the right plot illustrates that more than 60% of origin AS sets announce only a single MOAS prefix and around 90% announce fewer than ten MOAS prefixes. In addition to January 1, 2023, we also check one day per month and observe a similar pattern. Moreover, only a few origin AS sets announce large numbers of MOAS prefixes. Investigations show that this is again the TNet and Turk Telekomunikasyon merger case, see Section IV-A.

**To summarize:** The vast majority (95% for IPv4, 88% for IPv6) of MOAS prefixes are originated by two origin ASes, with some outlier prefixes being announced by more than 80 origins. Around 90% of origin AS sets announce fewer than ten MOAS prefixes. Therefore, the typical case is that two origin ASes announce a small number of MOAS prefixes.

## V. VISIBILITY

In this section we analyze the visibility of MOAS prefixes across route collector peers and investigate the minimum and

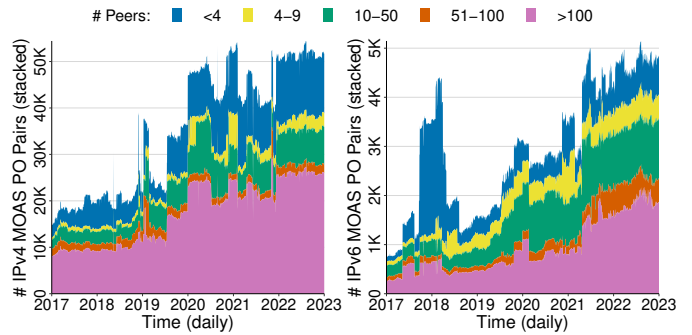


Fig. 9. MOAS prefix-origin pair visibility across route collector peers over time, for IPv4 (left) and IPv6 (right).

maximum visibility per prefix in detail.

### A. Visibility Across Route Collector Peers

First, we focus on the visibility of MOAS prefixes in route collector peers. RC peers are the ASes that feed routing information to the route collector via BGP sessions. Figure 9 illustrates the visibility of every MOAS prefix-origin (PO) pair, for IPv4 (left plot) and IPv6 (right plot), across RC peers for the period of the study. For IPv4, the largest group of MOAS PO pairs (i.e., around 50%), is visible in more than one hundred route collector peers (shown in pink), followed by MOAS PO pairs with a visibility of 3 or fewer peers (shown in blue). This “bimodal” visibility distribution is consistent with overall BGP prefixes [9]. PO pairs with a visibility between 10 and 50 peers (shown in green) form the third largest group.

In IPv6 we see a similar picture, with MOAS PO pairs with visibility  $>100$  being most common, followed by PO pairs with visibility between 10 and 50 peers. The  $<4$  category forms the third-largest group during our study period. However, from mid-2017 to mid-2018, we see a visible spike in IPv6 to around three thousand PO pairs in the  $<4$  category. We investigate this peak manually and find that it is largely due to hyper-specific MOAS prefixes originated by GTT Communications and DigitalOcean (see Section III-D) and /48 MOAS prefixes originated from the Thai eyeball network Realmove Company Limited (AS132061, AS7470).

**To summarize:** PO pairs of IPv4 MOAS exhibit a “bimodal” RC visibility behavior: Either they are widely visible in more than a hundred RC peers or very poorly visible in fewer than four RC peer. For IPv6 we see a similar picture, although semi-visible PO pairs are more common than in IPv4.

### B. Minimum and Maximum Visibility

To better understand the bimodal visibility behavior of MOAS prefixes, we investigate the minimum and maximum visibility (i.e., number of route collector peers) of every MOAS prefix depending on its origin AS. This allows us to see if a MOAS prefix has a similar visibility for all its origins or if the minimum and maximum are widely spread. In Figure 10 we show the minimum and maximum visibility of each PO pair for IPv4 (top) and IPv6 (bottom) for a single day (i.e.,

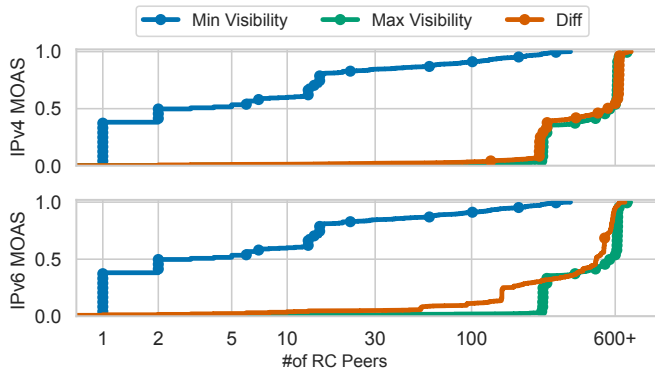


Fig. 10. Minimum, maximum, and difference visibility of MOAS prefixes on January 1, 2023 for IPv4 (top) and IPv6 (bottom).

January 1, 2023). We also check one day per month for the entire period of study, and observe a similar pattern.

For the maximum visibility we observe that for 99% of all MOAS prefixes—in both IPv4 and IPv6—at least one PO pair is visible by more than 100 RC peers. Moreover, for more than 50% of MOAS prefixes, at least one PO pair has a maximum visibility larger than 500 route collector peers. We attribute the steady increase in maximum visibility between 100 and 500 RC peers and after 500 RC peers to partial and full BGP feeds [42], respectively.

Analyzing the minimum visibility of MOAS prefixes shows that for around 40% of IPv4 and 20% of IPv6 MOAS prefixes, at least one PO pair is visible only at a single route collector peer. For over half of IPv4 MOAS prefixes, at least one PO pair has a minimum visibility of fewer than five RC peers, and more than 90% are visible on fewer than 100 RC peers. Around 70% of IPv6 MOAS prefixes have at least one PO pair with a minimum visibility of fewer than 100 RC peers.

This indicates, that each MOAS prefix is highly visible from one origin AS and has a very low visibility from another origin AS. To confirm this, we calculate the difference between the maximum and minimum visibility for each MOAS prefix. As can be seen in Figure 10, the difference distribution is very close to the maximum visibility, underlining that indeed for each MOAS prefix one origin is very visible whereas the other is barely visible.

**To summarize:** Each MOAS prefix has one origin, with which the prefix is very well visible at 100 or more RC peers and another origin which is barely visible at all. This hints at MOAS not being mainly used for anycast purposes, as for these we would expect a more balanced visibility distribution.

## VI. USERS AND USAGE OF MOAS PREFIXES

In this section we present an analysis on the users and usage of MOAS prefixes in the Internet. Specifically, we first examine the BGP relationships between the origin ASes of MOAS prefixes. Furthermore, we investigate the business relationships among MOAS prefix users and look into the use of MOAS prefixes by Hypergiants. Finally, we evaluate

to which extent MOAS prefixes are being utilized as anycast services.

### A. BGP Relationship of MOAS Origin ASes

First, we want to better understand the BGP relationship between origin ASes of MOAS prefixes. We use quarterly snapshots of CAIDA’s AS to organization [18] and AS classification inferences [16] datasets to identify the inter-AS BGP relationship of MOAS origin ASes. For simplicity, we focus on MOAS prefixes with two origins. These prefixes make up around 90% of all MOAS prefixes and two origin ASes simplify our analysis, as we do not have to consider ternary or higher order relationships.

We check for the following three categories of BGP relationships for MOAS prefixes. First, if both origin ASes have the same organization name in the dataset, we call them “Siblings”. Second, if any customer-to-provider or provider-to-customer relationship between the origin ASes of MOAS prefixes exists, we classify the inter-origin-AS relationship as “C2P/P2C”. Next, we examine if both origin ASes fall into the “Peering” category. Finally, if we do not observe any of the above known BGP relationships between the origin AS of a MOAS prefix, we mark the inter-origin-AS BGP relationship as a “No Relation Detected”.

Figure 11 illustrates the inter-origin-AS BGP relationship of MOAS prefixes for IPv4 (left) and IPv6 (right) over time. We find that almost half of all origin AS pairs for both IPv4 and IPv6 are classified as C2P/P2C (shown in blue). Peering relationships (shown in yellow) are less frequent, with siblings making up the smallest part of MOAS origin AS pairs. We do not detect any relationship using the CAIDA datasets for around half of all origin AS pairs. Overall, the distribution of categories remains relatively stable over time. Furthermore, when using the number of PO pairs as a metric (not shown), we see a similar trend for the share of relationship categories as in Figure 11, with a comparatively higher share of “Siblings” for two snapshots around 2020.

**To summarize:** The vast majority of detected MOAS origin AS relationships fall into the customer to provider category. This shows, that many MOAS prefixes are not related to sibling ASes, rather a provider announces the same prefix as the customer.

### B. Relationship and Business Type of MOAS Origin ASes

To shed more light on the business relationship between MOAS origin AS pairs, we now apply the 17 top-level (“layer 1”) categories of the ASdb dataset [22] on the most recent snapshot of MOAS prefix origin ASes (January 1, 2023). We also check MOAS origins in all available ASdb datasets for the past two years and observe a similar pattern.

As in the previous analysis, we focus on MOAS prefixes with exactly two origin ASes, i.e., 2618 unique origin ASes in total. For 97.2%, we find a matching business category in ASdb, out of which 77.6% match a single, and 23.4% origin ASes match more than one category. For easier understanding, we focus our analysis on ASes with one business category.




Fig. 11. BGP relationship of MOAS origin AS pairs for IPv4 (left) and IPv6 (right).

Figure 12 shows a heatmap of prevalent business categories for MOAS origin AS pairs. The number and color of every cell within the heatmap represents the prevalence of each category pair for MOAS origins. Examining the heatmap cells, we notice that in roughly 40% of the cases, both MOAS origins fall into the “IT” category. Furthermore, we notice a high number of IT company pairs with other business types, such as “Service”, “Retail”, or “Media”. This indicates that in addition to advertising their own prefix, organizations leverage IT companies to re-advertise prefixes, resulting in a MOAS event. Moreover, for every business type, we see an elevated level on the diagonal line, which means that ASes originate MOAS prefixes from the same business type. Overall, 89% of MOAS prefixes have at least one origin AS classified as an IT company. To exclude a possible bias towards IT companies, we analyze the prevalence of the IT category in the ASdb dataset. We find that MOAS origins from the IT category have a 10 percentage points higher share than general origin ASes in the ASdb dataset (56% vs. 46%).

**To summarize:** *MOAS prefixes are mostly originated by IT companies, with the same company type for both origins being most common.*

### C. Hypergiants Using MOAS Prefixes

Next, we investigate to what degree Hypergiants make use of MOAS prefixes. Hypergiants are large content providers, CDNs, and cloud networks (e.g., Google, Amazon, Meta) with a heavy-outbound traffic profile [43]. We use the list of 16 Hypergiants by Gigis et al. [44] and find that 11 of them (i.e., Akamai, Alibaba, Amazon, CDNetworks, Cloudflare, Facebook, Google, Incapsula, Limelight, Netflix, and Verizon) announce a total of 313 MOAS prefixes. The top three Hypergiants making use of MOAS prefixes are Verizon (98 MOAS prefixes), Netflix (48), and Google (45). This shows that MOAS prefixes are used by big players in the Internet, possibly to improve their network’s resilience, performance, and quality of experience for their users.

**To summarize:** *11 out of 16 Hypergiants use MOAS prefixes to reap the benefits of their users being able to access their networks via multiple origins.*




Fig. 12. Heatmap of business type relationships between MOAS origin pairs.

### D. Anycast MOAS Prefixes

MOAS prefixes can be good candidates for anycast services as they originate from multiple ASes. However, we find that only 225 (0.9%) of IPv4 and 89 (6.3%) of IPv6 MOAS prefixes are identified as anycast prefixes [24].

We observe that the number of origins per prefix for anycast MOAS prefixes is higher than for general MOAS prefixes: While around 90% of MOAS prefixes have two origin ASes, this percentage is less than 40% for anycast MOAS prefixes. More than 50% of anycast MOAS prefixes are announced by more than ten origin ASes. The prefix with the most origin ASes (61 origins) is the IPv6 prefix 2001:503:231d::/48, which is attributed to Verisign.

Even though the official root-server website lists single origin ASes for each of the 13 root DNS servers [45], we observe that the IPv4 address of A and J root DNS servers, operated by Verisign, use MOAS prefixes with a /24 CIDR size. In contrast to IPv4, we do not find any IPv6 MOAS prefix for DNS root servers.

**To summarize:** *The use of anycast within MOAS prefixes is low, with under 1% for IPv4 and around 6% for IPv6. When MOAS prefixes are used for anycast, more than 50% of them have more than ten origins. This shows, that MOAS prefixes are rarely used for anycast purposes, but if they are, they use a large number of origin ASes.*

## VII. RELATED WORK

In this section we present related work in the areas of prefix hijacking and MOAS prefixes.

**Prefix Hijacking:** When an attacker announces someone else’s prefix in BGP from its own AS, the prefix is being hijacked. A prefix hijack can cause service interruptions and consequently lead to traffic, financial, and reputation loss for the victim. The first well-known prefix hijacking case happened in 2008 when Pakistan Telecom hijacked the YouTube prefix for censorship



purposes [46]. Since then multiple prefix hijacking events have occurred. A more recent example of such an incident happened in August 2022 when several crypto services hosted on AWS were hijacked, resulting in financial loss [47]. Previous work has suggested various techniques and tools to detect and mitigate BGP prefix hijacks [6]–[11]. Websites such as BGPmon [4] and CAIDA’s Hijack Observatory [5] are monitoring the BGP for possible prefix hijack events. In 2018, Sermpezis et al. conducted a survey among 75 network operators which showed that operators mostly rely on third parties such as BGPmon [4] to detect prefix hijacks [12]. In 2019, Cho et al. presented a classification of BGP hijacks into different categories and found that 4%, 1%, and 2% of hijacks are due to typos, prepend mistakes, and BGP hijacking with a forged AS path, respectively [13].

**MOAS Prefixes:** In 2001, Zhao et al. investigated MOAS prefixes for the first time on the Internet using Routeviews route collectors [15] finding that the majority of MOAS conflicts were short-lived, lasting only a few days [30]. In 2007, Chin analyzed MOAS prefixes observed over 21 days and concludes that MOAS constitute a small but growing percentage of reachability information [31]. In 2011, Bornhauser et al. present MOAS Analyzer, a tool to identify, analyze, and automatically classify MOAS conflicts [32]. In 2014, Jacquemart et al. performed a longitudinal analysis of MOAS prefixes and showed that short-lived MOAS events are not due to misconfigurations, but because of origin instability or route flapping [33]. In 2015, Schlamp et al. investigated subprefix MOAS events using IRR data, topology data, and TLS certificates and found that the majority of these announcements with multiple origins are harmless [34].

To the best of our knowledge, this is the first work to perform an in-depth longitudinal analysis of long-lived MOAS prefixes.

## VIII. CONCLUSION

In this paper we analyzed long-lived MOAS prefixes over a period of six years. In IPv4 we identified around 24k long-lived MOAS prefixes compared to 1.4k in IPv6. We found that the vast majority of MOAS prefixes with entries in RPKI have a valid ROV state for all origins. Moreover, we found MOAS prefixes to use very specific CIDR sizes, indicating their use for fine-grained traffic steering. We identified that a considerable number of MOAS prefixes are related to mergers and acquisitions of companies. We recommend that network operators should clean up the extra MOAS produced due to mergers or other situations to prevent unnecessarily routing entries in the Internet routing environment. Furthermore, over 90% of MOAS prefixes were announced by two origin ASes, with the majority of origin pairs being in a customer-provider relationship. Finally, we have shown that the majority of MOAS users are IT companies and that MOAS are rarely used for anycast purposes. 🍌

## REFERENCES

- [1] Y. Rekhter (Ed.), T. Li (Ed.), and S. Hares (Ed.), *A Border Gateway Protocol 4 (BGP-4)*, RFC 4271 (Draft Standard), RFC, RFC Editor, Jan. 2006. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc4271.txt>.
- [2] J. Hawkinson and T. Bates, *Guidelines for creation, selection, and registration of an Autonomous System (AS)*, RFC 1930 (Best Current Practice), RFC, RFC Editor, Mar. 1996. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc1930.txt>.
- [3] S. Goldberg, “Why is it taking so long to secure internet routing?” *Communications of the ACM*, vol. 57, no. 10, 2014.
- [4] BGPmon, *BGPmon Hijacking events*, <https://www.bgpmmon.net/category/hijack/>.
- [5] CAIDA, *BGP Hijacks Observatory*, <https://dev.hicube.caida.org/feeds/hijacks/events>.
- [6] P. Sermpezis, V. Kotronis, P. Gigis, X. Dimitropoulos, D. Cicalese, A. King, and A. Dainotti, “Artemis: Neutralizing bgp hijacking within a minute,” *IEEE/ACM Transactions on Networking*, vol. 26, no. 6, 2018.
- [7] L. Qin, D. Li, R. Li, and K. Wang, “Themis: Accelerating the Detection of Route Origin Hijacking by Distinguishing Legitimate and Illegitimate MOAS,” in *USENIX Security*, 2022.
- [8] H. Imai, M. Okada, and A. Kanaoka, “Poster: MOAI: Multiple Origin ASes Identification for IP Prefix Hijacking and Mis-Origination,” in *NDSS*, 2020.
- [9] C. Testart, P. Richter, A. King, A. Dainotti, and D. Clark, “Profiling BGP serial hijackers: capturing persistent misbehavior in the global routing table,” in *ACM IMC*, 2019.
- [10] L. Oliver, G. Akiwate, M. Luckie, B. Du, and k. claffy kc, “Stop, DROP, and ROA: Effectiveness of Defenses through the lens of DROP,” in *ACM IMC*, 2022.
- [11] E. Biersack, Q. Jacquemart, F. Fischer, J. Fuchs, O. Thonnard, G. Theodoridis, D. Tzovaras, and P.-A. Vervier, “Visual analytics for BGP monitoring and prefix hijacking identification,” *IEEE Network*, vol. 26, no. 6, 2012.
- [12] P. Sermpezis, V. Kotronis, A. Dainotti, and X. Dimitropoulos, “A survey among network operators on BGP prefix hijacking,” *ACM SIGCOMM CCR*, vol. 48, no. 1, 2018.
- [13] S. Cho, R. Fontugne, K. Cho, A. Dainotti, and P. Gill, “BGP hijacking classification,” in *IEEE TMA*, 2019.
- [14] RIPE NCC, *RIPE Routing Information Service*, <http://www.ripe.net/ris/>, 2022.
- [15] University of Oregon, *Route Views Project*, <http://www.routeviews.org/>, 2022.
- [16] CAIDA, *AS Relationships Dataset*, <https://www.caida.org/catalog/datasets/as-relationships/>.
- [17] M. Luckie, B. Huffaker, A. Dhamdhere, V. Giotsas, and K. Claffy, “AS relationships, customer cones, and validation,” in *ACM IMC*, 2013.
- [18] CAIDA, *Inferred AS to Organization Mapping Dataset*, <https://www.caida.org/catalog/datasets/as-organizations/>.
- [19] CAIDA, *Mapping Autonomous Systems to Organizations: CAIDA’s Inference Methodology*, <https://www.caida.org/archive/as2org/>.
- [20] Z. Chen, Z. S. Bischof, C. Testart, and A. Dainotti, “Improving the Inference of Sibling Autonomous Systems,” in *PAM*, 2023.
- [21] A. Arturi, E. Carisimo, and F. E. Bustamante, “as2org+: Enriching AS-to-Organization Mappings with PeeringDB,” in *PAM*, 2023.
- [22] M. Ziv, L. Izhikevich, K. Ruth, K. Izhikevich, and Z. Durumeric, “ASdb: a system for classifying owners of autonomous systems,” in *ACM IMC*, 2021.
- [23] RIPE NCC, *RPKI snapshots*, <https://ftp.ripe.net/rpki/>.
- [24] BGP.tools, *Anycast Prefixes list on GitHub*, <https://github.com/bgptools/anycast-prefixes>, 2022.
- [25] R. Sommesse, L. Bertholdo, G. Akiwate, M. Jonker, R. van Rijswijk-Deij, A. Dainotti, K. Claffy, and A. Sperotto, “Manycast<sup>2</sup>: Using anycast to measure anycast,” in *ACM IMC*, 2020.
- [26] J. Mitchell, *Autonomous System (AS) Reservation for Private Use*, RFC 6996 (Best Current Practice), RFC, RFC Editor, Jul. 2013. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc6996.txt>.
- [27] IANA, *IANA IPv4 Special-Purpose Address Registry*, <https://www.iana.org/assignments/iana-ipv4-special-registry/iana-ipv4-special-registry.xhtml>, 2023.

- [28] IANA, *IANA IPv6 Special-Purpose Address Registry*, <https://www.iana.org/assignments/iana-ipv6-special-registry/iana-ipv6-special-registry.xhtml>, 2023.
- [29] M. Cotton, L. Vegoda, and D. Meyer, *IANA Guidelines for IPv4 Multicast Address Assignments*, RFC 5771 (Best Current Practice), RFC, RFC Editor, Mar. 2010. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc5771.txt>.
- [30] 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 *ACM IMW*, 2001.
- [31] K.-W. Chin, "On the characteristics of BGP multiple origin AS conflicts," in *IEEE ATNAC*, 2007.
- [32] U. Bornhauser, J. Luehr, M. Schaefer, T. Trimborn, and Z. U. Huda, "Automatic Analysis and Classification of Multiple Origin AS (MOAS) Conflicts," in *IEEE LCN*, 2011.
- [33] Q. Jacquemart, G. Urvoy-Keller, and E. Biersack, "A longitudinal study of BGP MOAS prefixes," in *IEEE TMA*, 2014.
- [34] J. Schlamp, R. Holz, O. Gasser, A. Korsten, Q. Jacquemart, G. Carle, and E. W. Biersack, "Investigating the Nature of Routing Anomalies: Closing in on Subprefix Hijacking Attacks," in *IEEE TMA*, 2015.
- [35] V. Satopaa, J. Albrecht, D. Irwin, and B. Raghavan, "Finding a 'Kneedle' in a Haystack: Detecting Knee Points in System Behavior," in *IEEE ICDCS*, 2011.
- [36] K. Z. Sediqi, L. Prehn, and O. Gasser, "Hyper-Specific Prefixes: Gotta Enjoy the Little Things in Interdomain Routing," *ACM SIGCOMM CCR*, vol. 52, 2 Jun. 2022.
- [37] NIST, *NIST RPKI Monitor*, <https://rpk-monitor.antd.nist.gov/>, 2023.
- [38] Daily Sabah, *Türk Telekom swallows up Avea and TTNET*, <https://www.dailysabah.com/business/2016/01/27/turk-telekom-swallows-up-avea-and-ttnet>.
- [39] L. Abboud and J. Toyer, *Orange strikes 3.4 billion euro deal to buy Spain's Jazztel*, <https://www.reuters.com/article/us-orange-jazztel-idUSKBN0HA1SM20140915>, Sep. 2014.
- [40] GTT, *GTT Completes Acquisition of KPN International*, <https://www.gtt.net/us-en/resources/news/gtt-completes-acquisition-of-kpn-international>, Dec. 2019.
- [41] IPv4 and IPv6 reports, *AS65000 and AS131072 BGP Routing Table Analysis Report*, <https://bgp.potaroo.net/as2.0/bgp-active.html>.
- [42] V. Giotsas, M. Luckie, B. Huffaker, and K. Claffy, "Inferring complex AS relationships," in *ACM IMC*, 2014.
- [43] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, "Internet inter-domain traffic," *ACM SIGCOMM CCR*, vol. 40, no. 4, 2010.
- [44] P. Gigis, M. Calder, L. Manassakis, G. Nomikos, V. Kotronis, X. Dimitropoulos, E. Katz-Bassett, and G. Smaragdakis, "Seven years in the life of Hypergiants' off-nets," in *ACM SIGCOMM*, 2021.
- [45] IANA, *root-servers.org*, <https://root-servers.org/>, 2023.
- [46] RIPE NCC, *YouTube Hijacking: A RIPE NCC RIS case study*, <https://www.ripe.net/publications/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study>, Mar. 2008.
- [47] D. Goodin, *How 3 hours of inaction from Amazon cost cryptocurrency holders \$235,000*, <https://arstechnica.com/information-technology/2022/09/how-3-hours-of-inaction-from-amazon-cost-cryptocurrency-holders-235000/>, Sep. 2022.