Serial BGP Hijackers: A Reproducibility Study and Assessment of Current Dynamics

Ebrima Jaw[†], Moritz Müller[†*], Cristian Hesselman[†*], Lambert Nieuwenhuis[†]
[†]University of Twente, Enschede, The Netherlands
[*]SIDN Labs, Arnhem, The Netherlands

***Abstract*—The Border Gateway Protocol (BGP) is the Internet's most crucial protocol for efficient global connectivity and traffic routing. However, BGP is well-known to be susceptible to route hijacks and leaks. Route hijacks are the illegitimate announcements of network resources, intentionally or unintentionally, which can compromise the confidentiality, integrity, and availability of communication systems. In the past, so-called "serial hijackers" have hijacked Internet resources multiple times, some lasting for several months or years. So far, only the paper "Profiling BGP Serial Hijackers" focuses explicitly on those repeated offenders, and their study dates back to 2019. Back then, they had to process large amounts of BGP announcements to find a few potential serial hijackers. In this paper, we revisit the profiling of serial hijackers. We reproduce and extend the study from 2019 and show that we can identify potential offenders with less data while achieving similar accuracy. We show that most of the alleged serial hijackers are still active on the Internet, announcing prefixes that belong to other ASes. In conclusion, our study confirms that there has been no significant increase in the evolution of serial hijacking activities during the last five years. However, we found that the active alleged serial hijackers and the identified potential malicious actors still threaten the Internet's security and stability.**

I. Introduction

The Border Gateway Protocol (BGP) is crucial to the Internet's scalability and resilience. It enables individual networks, so-called autonomous systems (ASes), to form a network of networks and to communicate without centralized control [1]. This decentralized control approach has enabled the growth of the Internet but also made it more vulnerable to misconfigurations and malicious actors [2].

Route or prefix hijacks are events in which an AS intentionally or unintentionally originates reachability information of another network [3], negatively impacting Internet services frequently. For instance, in 2020, Rostelecom (AS12389) hijacked over 8,000 prefixes belonging to more than 200 cloud providers and CDNs, such as Cloudflare, Akamai, Amazon, Google, and Facebook that lasted for almost an hour [4]. The research community and industry have documented many more incidents [1], [5]–[7].

Despite efforts to prevent and mitigate such attacks [1], [7], hijacks are still a threat [8]. Most of the time, offenders hijack one AS or prefix at once. However, Testart et al. have shown in 2019 that occasionally, networks perform multiple hijacks repeatedly and for several days, weeks, or even months [9]. Their study found around 900 of such potential *serial hijackers* active between 2014 and 2018.

978-3-903176-64-5

Since then, serial hijackers have not received any attention from the research community anymore. For this reason, the method by Testart et al. remains unreproduced; it remains unclear to what extent serial hijackers are still a problem and whether serial hijackers identified by Testart et al. are still active. Our study attempts to fill this gap and makes the following contributions:

- We validate findings by the original paper, demonstrating that the observations by Testart et al. still largely hold.
- While the previous study relied on large data sets that map prefixes to their announcing ASes with a five-minute granularity, we show that we can characterize and identify serial hijackers with 50 times less data, thereby showing that studying (serial) hijackers is also feasible for researchers with fewer resources.
- We validate the results of the existing classifier used in the original study, confirm their reported performance, and develop our competitive classifier relying on fewer features.
- We show how serial hijacking activities have evolved in the last four years, contributing to discussions on mitigating BGP route hijacking in general and serial hijackers in particular.

The remaining paper is structured as follows: We summarized background information in Section II. Sections III and IV present an overview of serial hijackers and our datasets. Sections V to VIII present our findings, and Sections X, IX, and XI present our discussion, related work, and conclusions.

II. Background

This section summarizes the fundamental concepts of BGP, its security vulnerabilities, and the currently available and emerging countermeasures.

A. Core Functionalities

Routers at the edge of ASes use BGP to exchange reachability information about IP address space with other ASes' routers. These routers calculate the best route to reach an IP prefix based on the received information. The definition of the best route depends on the connected ASes' routing policies and is influenced by organizational and business relationships [7]. Routers generally prefer routes that claim a path towards a more specific part of a prefix (e.g., preferring a path towards

192.0.2.0/24 over a path towards 192.0.2.0/20), that provide connectivity through a customer or peer over a route through a transit provider, or that provide a shorter path towards a prefix (fewer hops involved in routing the traffic).

B. BGP Hijacks

BGP lacks built-in authentication and validation mechanisms for connected ASes to verify the authenticity of available routes. These vulnerabilities regularly cause wide-scale malicious or unintentional hijacks, compromising access to critical systems and services [1], [6], [9]. The two main classifications of BGP hijacks are prefix and AS path hijacking: Prefix hijacking is an illicit origination of IP prefixes or blocks belonging to another AS(es) [10]. This type of hijacking event will cause multi-origin AS (MOAS) events, which means that prefixes originate from both the victim and the malicious network [11]. In a path hijack, malicious actors can manipulate the path by claiming that the attacker's AS lies on the path towards the victim's AS. This route will likely impact a wide range of ASes if advertised as more specific.

C. Countermeasures

Even though they were already available during the previous study in 2018, the adoption of technologies to partially prevent hijacks has grown considerably. For example, route origin validation, based on the Resource Public Key Infrastructure (RPKI), has increased recently (from 6% to 12% between 2022 and 2023 [12]). However, route origin validation only protects against origin hijacks, and the question remains to what extent it has contributed to decreasing the number of hijacks in general and serial hijackers specifically. Other measures like Autonomous System Provider Authorization (ASPA) [13] and BGPSec [14] are either in their early deployment stage or have yet to be deployed in the wild. Aside from Testart et al. several studies have proposed methods to detect standalone hijacks, and we will discuss the most relevant attempts in Section IX.

III. Serial Hijackers

Testart et al. [9] coined the term *serial hijackers* to refer to malicious actors that repeatedly carry out *origin hijacks* on various networks, some lasting for months or years. Although their work remains the only longitudinal study analyzing the behavior of these repeated offenders, we have found reports in the NANOG mailing list (2014-2018) anecdotally supporting their claims.[1]

A. Profiling Serial Hijackers

Testart et al. selected a list of ASes from the MANRS initiative [15] and other large providers to create a dataset of 217 legitimate ASes. Additionally, they processed five years of emails sent on the NANOG mailing list to extract 23 frequently reported malicious ASes that have hijacked many address blocks for an extended period.

In order to characterize and subsequently identify serial hijackers, they relied on 3 data sets: *First*, they used CAIDA's BGPview tool [16] to create a longitudinal dataset of mappings between IP prefix and announcing AS number. This data consisted of five-minute snapshots of the routing table, resulting in 525,888 snapshot files across five years (2014-2018). The authors used this data to identify characteristics that distinguish serial hijackers from "regular" ASes, namely: (i) serial hijackers do not announce prefixes continuously, (ii) the number of announced prefixes fluctuates, (iii) and prefixes are announced for a short time. *Second*, they derived a list of MOAS conflicts from the same source, showing that serial hijackers announce MOAS-conflicts more often. *Third*, they relied on RIR delegation files to identify which prefix is allocated by which RIR [17]. They found that serial hijackers announce prefixes allocated by different RIRs more often than regular ASes.

Testart et al. transformed these characteristics into statistical features used to train a classifier relying on ExtraTrees. This classifier flagged about 900 ASes as potential serial hijacker. Finally, they used other supplementary data to cross-validate the flagged networks, such as SPAMHAUS ASN-DROP List [18] and UCE-PROTECT Level 2 spam blacklist [19]. We refer the reader for a detailed description to their paper.

B. Reproducibility: Our Motivation

To our knowledge, Testart et al. are the only ones who studied the behavior of serial hijackers [9], which dates back to 2019. Our goal is to study the development of serial hijackers since then. However, when trying to reproduce their methodology, we found that the authors relied on terabytes of BGP data, which requires storage capacities and processing capabilities that we did not have.

The increase of data collected at route collectors does not only pose a challenge for us [20]. A recent study by Thomas et al. highlighted the burden of processing these potentially redundant BGP updates and proposed an approach to mitigate these challenges [21]. Therefore, the motivation of our research is not only to validate the method by Testart et al. and apply it to recent data but also to reproduce their work with less data. With this, we also would like to enable other researchers with similar constraints as ours to study this relevant topic. In the end, we want to answer the following research questions:

1) Can we reproduce the findings of the original work using less data?
2) How did serial hijacking events evolve since 2019?
3) What actions did the Internet community take in response to the reported alleged serial hijackers?

IV. Datasets

a) Longitudinal BGP dataset: We found that the five-minute snapshots of the peer-pfx-origins dataset from CAIDA are highly redundant. For this reason, we select only a single five-minute snapshot per day of the peer-pfx-origins dataset, and we show in Section V that this is sufficient to characterize the attributes of regular and malicious networks.

[1] e.g. https://mailman.nanog.org/pipermail/nanog/2018-August/096737.html

Since the peer-pfx-origins dataset is only available until the end of 2018, we use CAIDA's BGPView tool [16] to compute our current five-minute snapshots of the peer-pfx-origins dataset for 2019-2023 using BGP updates from RIPE RIS [22] and Routeviews [23].

TABLE I: Dataset properties: Aggregates of the longitudinal BGP and supplementary datasets.

Details of longitudinal BGP Dataset				
	(2014 - 2018)		(2019 - 2023)	
	IPv4	IPv6	IPv4	IPv6
Start Date	01/01/14 00:00 UTC		01/01/19 00:00 UTC	
End Date	31/12/18 23:55 UTC		23/11/23 23:55 UTC	
Snapshot files	1,825	1,825	1,787	1,787
Unique prefixes	2,750,800	383,806	3,108,234	1,613,626
Unique origins	76,008	22,349	93,583	41,966
Prefix-origin pairs	3,245,325	410,454	3,590,425	1,742,951
Details of MOAS Dataset				
MOAS prefixes	269,809	16,814	194,773	32,483
MOAS origins	43,205	8,038	34,731	13,551
Details of RIR Dataset				
Start Date: 01/01/2019		End Date: 23/11/2023		
	IPv4	IPv6	ASN	
RIR Resources	304,159	387,203	126,758	

b) MOAS dataset: We create a MOAS dataset using ten years of longitudinal BGP dataset (2014-2023). We parse each five-minute snapshot to extract the list of ASNs originating the same prefixes simultaneously, as well as the unique originated prefixes and the counts of ASNs.

c) Processing: We follow the same data-preprocessing steps described in [9], such as discarding incomplete or corrupted snapshot files. However, considering we only use a single snapshot file, we do not discard prefixes with low visibility. We collect our snapshots every day at 00:00:00 UTC. We verified that our results were independent of the chosen time of the day by comparing the results of different single snapshot timestamps within a day. Each snapshot lists the observed prefixed, their announcing ASNs, and the number of peers that observe the prefix-origin pair.

Then, we adopted the same method as the original study to normalize our prefix visibility within the range of (0,1]. A prefix-origin pair with a visibility of 1 would be observed by all peers feeding their updates during that snapshot [9]. We grouped all prefix-origin pairs of each ASN into continuous timelines, which summarizes the long-term behavior of all the ASes in our dataset.

d) RIR delegation files: Delegation files are publicly available [17]. These files document the allocation of Internet resources such as IPv4, IPv6, and ASNs. We downloaded a single snapshot of these files at the start and end of every year across the five years for each RIR to calculate the RIR-based features. We adopted this approach due to the redundancy in processing all the daily snapshots across our study period. For instance, we mainly observed 100% duplicates of records between two subsequent snapshot files, such as the 13th and 14th of Nov. 2019, and the same applies to our weekly and monthly comparison.

Using Jaccard similarity, we compared the overlap of resources between the first snapshot of the year (1st Jan. 2019) and the snapshots of the rest of the year. We found that the snapshot of 1st Jan. has an average similarity of 0.941 with the remaining snapshots (minimum 0.924, maximum 0.960, and standard deviation of 0.011). We observed a difference of 0.0176 when comparing the first and last snapshots of the year, and this is due to the new records appearing in the latter snapshots. Despite these differences, we believe that using a subset of RIR delegation snapshots is sufficient for computing our features.

Table I presents the details of our reproducibility dataset for 2014-2018, the recently collected dataset (2019-2023), and the aggregated MOAS dataset, which will be available upon request due to the file size.

e) Ground-truth data: Finally, we utilized the same ground-truth dataset from the original study for consistency and validity, which consists of 23 serial hijackers and 217 legitimate ASes. We also adopted the metric of the original study to select 10,623 ASes that at least announced ten prefixes as our prediction set.

V. Characterizing serial hijackers with sparse data

This section validates our hypothesis that we can characterize the attributes of legitimate networks and serial hijackers with our sparser dataset.

A. Dominant characteristics of legitimate networks and serial hijackers

To test our hypothesis that using less data can capture the characteristics of serial hijackers, we compare the behavior described by Testart et al. with our observations. We find that the most highlighted characteristics in the original paper hold while using a significantly sparser dataset.

1) Prefix origination patterns: We show in Figure 1 which prefixes the ASes have announced over 5 years and how many peers have observed the announcement, but only using one five-minute daily snapshot. Each row represents one prefix. Figure 1a shows the colored normalized visibility of prefix announcement patterns of a legitimate AS, while Figure 1b shows the pattern of a serial hijacker [9].

In Figure 1a (and as shown by the original paper), we can see that British Telecom (AS5400) shows a relatively stable origination behavior during the five years, such as consistent prefix announcement duration (33.68% of active prefixes for the entire five year period) and moderately few prefixes compared to serial hijackers (≈200 and 800, respectively). Similar to the original study, we also observed temporal unstable behavior for regular ASes, such as no origination of prefixes (white spaces) for specific periods and short announcement lifetime of prefixes, which requires consideration for building reliable and effective classifiers [9].

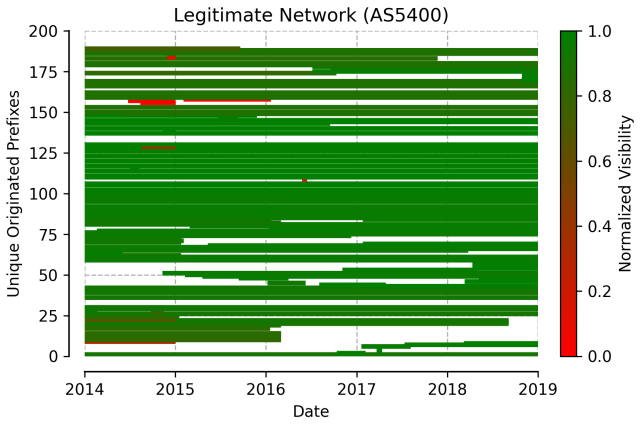

(a) Prefix origination pattern of British Telecom (AS5400).

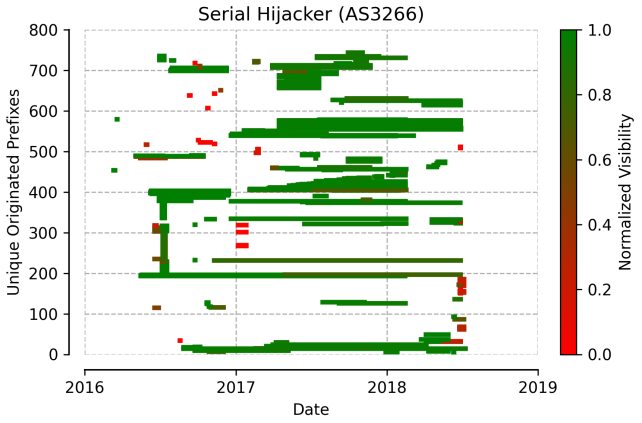

(b) Prefix origination pattern of a serial hijacker (AS3266).

Fig. 1: Reproduced five and three years prefix origination patterns of a legitimate AS (5400) and serial hijacking AS (197426) [9]

In contrast, AS3266 shown in Figure 1b depicts inconsistent origination patterns of a serial hijacker over three years, with most prefixes announced for a very short time. For instance, only 52.35% of the prefixes were active for only a year (2017-2018), and 0% of consistent active prefixes for the whole five-year period. Moreover, AS3266 showed erratic behavioral patterns on the global routing table, with about 800 announcements of prefixes within only three years, which is about 75% more prefixes than AS5400 for five years.

As with other hijacks, it is hard to say with certainty if a serial hijacker is acting with malicious intentions or whether misconfigurations cause the hijacks. However, the original study's claims that a single actor sometimes hijacks hundreds or thousands of prefixes over several years strongly indicate malicious intent, and we follow this reasoning in our study.

2) Prefix origination changes: Figure 2a shows a consistent origination behavior of AS7922 over five years. Our study missed the spike in Figure 2a in the original study. We confirmed this was due to the deaggregating of larger prefixes and was not related to hijacking.

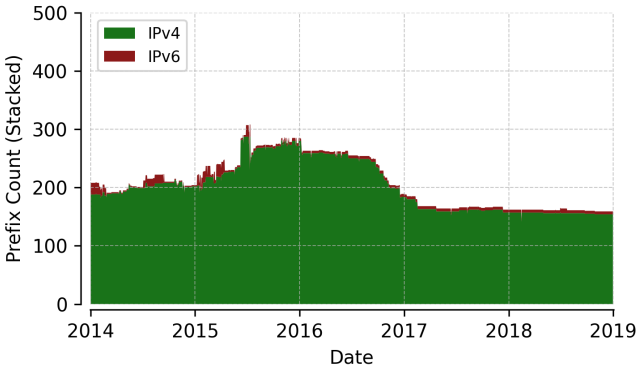

(a) Prefix origination changes for a normal AS (7922).

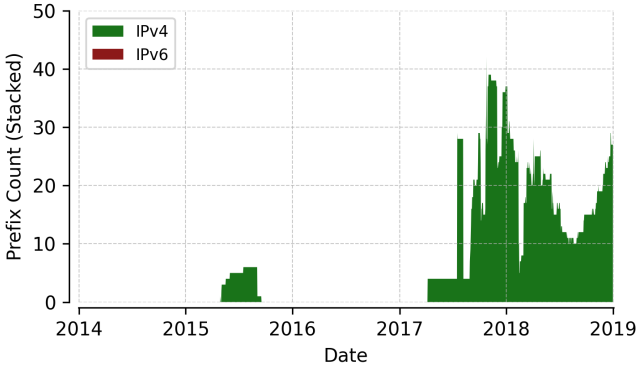

(b) Prefix origination changes for a malicious AS (133955).

Fig. 2: Reproduced prefix origination changes for a normal (AS7922) and malicious (AS133955) networks over the period of 5 years.

In contrast, Figure 2b shows the volatile and inconsistent origination pattern of malicious AS (133955). We observed that AS133955 had no BGP activity until mid-2015 when it first had few activities and was absent until after 2017. Subsequently, it recorded a high volume of activities until the end of 2019.

3) Longevity of the fraction of prefix-origin pairs over time: The duration of prefix announcements can also help us to characterize the serial hijackers from normal ASes. Figure 3a shows the fractions of prefix-origin pairs for AS57129, and we can see that about 90% of its IPv4 prefixes were announced between 0 and 250 days, and about 60% of its IPV6 prefixes were announced for less than 400 days. This behavior correlates with the attributes of a serial hijacker discussed in the previous study. In contrast, Figure 3b shows the prefix origination duration for AS7922, which shows that about 40% of its prefix announcements lasted between 500 and 2880 days. However, we can see that about 40% of the prefixes for AS7922 were announced for less than 50 days. This behavior happens due to local engineering practices we have investigated and validated as explained in (Section V-A2).

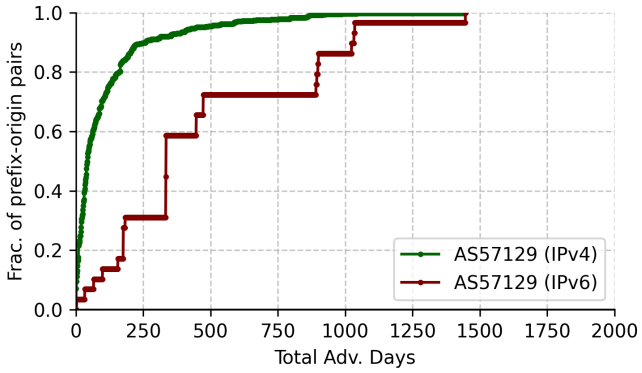

(a) Serial Hijacker (AS57129)

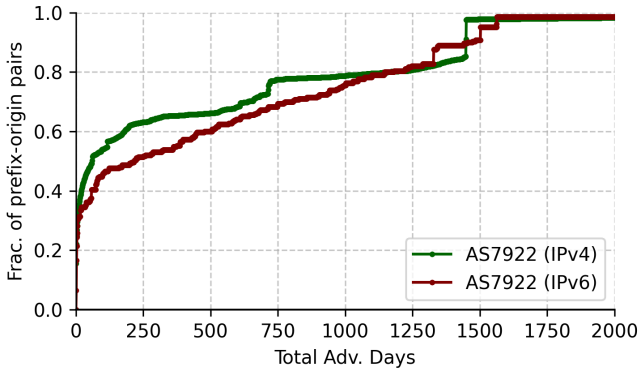

(b) Legitimate network (AS7922)

Fig. 3: Fraction of IPv4 and IPv6 prefix-origin pairs over the total advertisement days as announced by a serial hijacker (AS57129) and a legitimate network (AS7922)

VI. Detecting Serial Hijackers with Sparse Data

After showing that we can characterize serial hijackers with less data, we now want to *detect* serial hijackers automatically.

A. Assessment of the Original Classifier on historical data

We first attempt to validate whether the classifier provided by Testart et al. performs as well on sampled data as on the full data set. To do so, we rely on computed features from our sparse dataset for the ground truth (240 ASNs) and the prediction set (10,623 ASNs).[2] Then, we train and evaluate it using our sparse dataset.

Figure 4 illustrates the Mean Out-of-Bag (OOB) error of the original classifier on our sampled data. As in the previous study, we apply different sampling methods to cope with the unbalanced training set. We observed that increasing the forest size reduces mean OOB error for all sampling methods (and thereby improves OOB accuracy). Finally, our correlation analysis shows that some features were redundant, making no difference in the final prediction. For example, in most cases, the *weekly* and *monthly bins* features correlate with the *first* and *third-quartile advertisement time* and other BGP-based related features. As a result, we only used 44 features to train and evaluate our models.

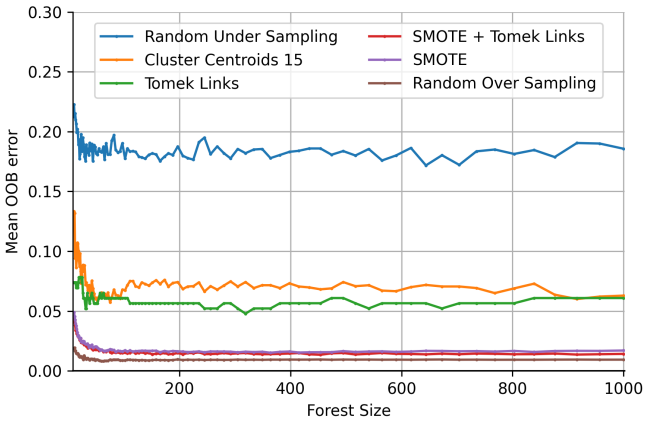

Fig. 4: The mean OOB error of different sampling methods and forest size of the original machine learning classifier on our sparse dataset.

B. Improving on the original classifier

Even though the original classifier showed good performance on full and sampled data (80% precision), Testart et al. already discussed that not all flagged ASes by the classifier were necessarily malicious. For this reason, we evaluate in this section if we can improve the classifier further. Selecting an effective model depends on the specific use case. Since accusing an AS of serial hijacking could have serious consequences, detecting serial hijackers with high confidence is crucial. We tested a Logistic Regression model, a Random Forest classifier, and an ExtraTree classifier to determine an effective and suitable model for our use case. In addition to the best-performing sampling methods in Figure 4, we also implemented the Proximity Weighted Random Affine Shadowsampling (ProWRAS) and Adaptive Synthetic Sampling (ADASYN) oversampling techniques to address the class imbalances within our ground truth dataset.

Our results on various metrics and oversampling methods with up to 1000 forest sizes show no significant improvement in our model performances after 200 trees. Table 2 presents the findings on the original published model[3] and ours with sampled current data (2019-2023).

TABLE II: Performances of the two classifiers

Models	Precision	Recall	Accuracy	F1-Score
Extra Trees (Reproduced)	0.88	0.99	-	-
Extra Trees (Ours)	0.97	0.75	0.94	0.82

[2]Before we could run their code, we upgraded the classifier (ExtraTrees) to Python 3. Additionally, we addressed some incompatibility issues in the code due to new updates in the `scikit-learn` library to align with the current Python standards and ensure continued functionality.

[3]We only reported the precision and recall as done in the original study

In the end, we opted for the Extra Trees classifier because it achieved the fewest false positives and negatives. Subsequent sections investigate the evolution of serial hijackers using this classifier.

VII. Serial hijackers 2019 - 2023

We applied our classifier on 10,622 ASes that originated at least 10 IPv4 prefixes from 2019-2023. It classified 766 ASes to have attributes similar to those of a serial hijacker, and we called them *flagged networks* henceforth. The subsequent sections present our findings on these flagged ASes.

A. Common flagged networks

We first compared our flagged ASes to those by Testart et al. in 2019. We found that both models have flagged 279 ASes (36.41% of our flagged networks) as having the characteristics of a serial hijacker, confirming a reasonable agreement among both models on the existence of these potentially malicious networks. It is evident that there is a difference in the remaining classified networks, possibly due to our sampling methods, and we have computed various metrics presented in Table 2 that illustrate the detailed performance assessment of our final model.

B. BGP misconfigurations

A small fraction of the flagged ASes could have resulted from BGP misconfigurations, such as the leaking of private ASNs and the fat finger errors due to path prepending [9]. We have found 14 private and 5 single-digit networks within our flagged ASes ($\approx$ 1.83% and 0.7% of the flagged networks), respectively. We found evidence of such incidents while processing the NANOG mailing list. For instance, Iranian state telecom TIC used a private ASN to hijack a site hosting adult content, which unfortunately leaks to the global routing table [24] and reports of single-digit ASes like AS6 and AS2 announcing prefixes belonging to other networks[4]. We excluded the private and single-digit networks from our flagged networks and used the remaining 747 flagged networks for our subsequent analysis.

C. Benign serial hijackers

DDoS protection networks like Path Network (AS396998) exhibit the exact characteristics of serial hijackers, and our dataset contains a significant share of MOAS prefixes originated by some DDoS protection networks, such as Vercara (Neustar). We compared our flagged networks to a manually compiled list of 70 DDoS protection services and found our classifier to have flagged 14 ASes (1.8817% of our flagged networks). The manually compiled list is small, and we could miss other networks.

[4]https://seclists.org/nanog/2018/Apr/206

D. Listed flagged ASes

We used SPAMHAUS's Don't Route or Peer Drop List [18], Clean Talk-listed ASNs [25], and UCEPROTECT's Level 3 list [19] to determine potential malicious activities for our remaining 733 flagged networks. Our classifier flagged 45 (6.53%) ASes from the above lists. Clean Talk keeps track of lists of ASes with active spam IP addresses, and we found 20 of our flagged networks to be on this list. Additionally, we found 17 of our flagged networks on the UCEPROTECT's Level 3 list of ASes and 8 ASes in the SPAMHAUS's Don't Route or Peer Drop List. Although the share of our flagged ASes in the above lists is less compared to the claims of the original study, we can still observe some potential malicious activities among our flagged networks.

E. Sign of spamming among flagged ASes

We used 1360 prefixes from Spamhaus DROP, DROPv6, and extended Drop (eDROP) across the duration of our study to determine spamming activities among our flagged ASes. These lists contain RIR allocations and sub-allocated blocks of hijacked prefixes or malicious netblocks [18]. We mapped the timestamps of our daily BGP snapshots to the timestamps of Spamhaus's lists. Then we checked if any of our flagged ASes announced any listed prefixes in matching timestamps and returned the prefixes and the flagged ASes announcing them. Out of the remaining 688 flagged ASes, we observed that 84 (12.21%) of our flagged ASes have simultaneously announced 1230 (90.44%) of the listed prefixes across various timestamps. These consistent announcements of blacklisted prefixes indicate spamming activities, thereby making this group potentially malicious ASes.

F. Effect of MOAS

The original paper uses about 7 MOAS features to differentiate between legitimate and serial hijackers. Using MOAS features to classify malicious announcement activities is a logical decision. However, while creating our MOAS dataset, we found a significant increase in the use of MOAS for legitimate use cases. For instance, major networks like Vercara (Neustar) and Verisign were mainly responsible for an increase of about 1.02% of MOAS between 2020 and 2023.

Overall, we found that 508 of our flagged ASes have announced MOAS since 2019; this is 22.5% times more often than the AS not flagged by our classifier. These erratic origination patterns of the flagged ASes match the characteristics of serial hijackers, making our model classify these ASes as malicious. Also, potential hijackers flagged by the previous study announced 6.25 times more MOAS. We manually inspected the MOAS-flagged ASes and found that many flagged ASes provide security solutions. This indicates that we need to better understand the motivation behind announcing MOAS to improve the classifier [26].

G. Fate of the remaining flagged ASes

We used CAIDA's ASRank API [27] to determine the status of our remaining flagged ASes. Our findings show that

there are only 4 ASes (AS49121, AS49121, AS264009, and AS263481) with at least one peer out of the remaining 96 flagged ASes, while 91.94% of the ASes have no peering relationship. About 50% of these ASes currently have no provider. Our findings confirmed that they have been disappearing from the global routing table between 2019 and 2024. Although we do not know why these ASes went out of operation, it is uncommon for a stable network to disappear entirely from the global routing table.

VIII. The fate of serial hijackers

Testart et al. relied on 23 ASes accused of serial hijacking on public mailing lists. In this section, we discuss what happened to these ASes after they have been called out and after Testart et al. published their study. Here, we rely on data from RIPE RIS [28] and archives containing information about IP address space allocations and assignments of the five regional RIRs [17].

Of the 23 ASes, 9 (39.0%) were still active by the end of 2022. Their AS number was continuously allocated from when they were accused of hijacking Internet resources and were still announcing prefixes. The second largest group is ASes, which were not allocated for some period after their hijacking activities but were reallocated by the RIRs at some point. This group consists of 8 ASes (34.8%). It is unclear whether the new owner of the AS is a different entity than the one carrying out the hijacks. 4 ASes (17.4%) are not allocated anymore and were not reallocated in the meantime, and 2 ASes (8.7%) are still allocated to the same entity but do not announce any prefixes visible in RIS. We refer to the first group of ASes as *active ASes* and to the remaining groups as *passive ASes*.

Active ASes. Active ASes announce more prefixes today than before 2019. The median number of daily announced prefixes rose from 30 to 131, with one AS announcing a maximum of 2,838 prefixes. Even up to date, we find that these ASes are occasionally or continuously announcing MOAS. At some point, one AS announced almost 20 times more MOAS than its peak before 2019. However, we could not verify whether these announcements were actual hijacks.

Passive ASes. ASes of serial hijackers are more often released (not allocated in the RIRs' databases) than other ASes. 12 ASes (52,2%) were unallocated at some point after the hijacking events. In comparison, only 10.4% of the RIPE ASes active between 2019 and 2023 were unallocated at some point.

The time it takes until a hijacking AS is de-allocated varies. ASes stayed assigned for a median of 239 days after their last BGP announcement, but one AS was still allocated after 1,609 days. In contrast, two ASes announced prefixes even though their AS was no longer allocated for 20 and 24 days. Announcements by these two ASes should be considered invalid. Of the ASes that were de-allocated at some point in time after the hijacks, all belonged to the RIPE NCC RIR. It remains unclear whether the de-allocation was due to actions taken by this RIR or some other factor.

Unclear fate. The fact that 60% of the alleged hijackers are no longer active indicates that their behavior was at least irregular. However, whether these ASes became nonactive voluntarily or due to actions like de-peering or de-allocation is unclear. At the same time, almost 40% of the alleged hijackers are more active than ever on the global routing table. One possible explanation could be that their irregular behavior was due to mismanagement and has been resolved since then. Another explanation could be that the community did not punish them for their malicious behavior.

IX. Related Work

The research community has introduced different approaches to detect and classify routing hijacks. We briefly summarize a few selected works here.

Xiang et al. [29] proposed Argus, a hijacking detection system that uses control and data plane information to identify malicious routes, validated using active routable IP addresses. Argus monitored ≈12,000 anomalies and detected 63 possible hijacking incidents.

Sermpezis et al. [6] leveraged various datasets to propose an automated prefix hijacking detection and mitigation system (ARTEMIS). It enables network operators to run the system based on their internal configuration requirements to maintain privacy and improve real-time detection. However, ARTEMIS cannot detect attacks intended for other prefixes excluded within the network and observes delays in accessing live BGP feed that could affect performance.

Cho et al. [30] applied the concept of AS hegemony with heuristic methods to classify BGP hijacks into pre-pending mistakes, origin changes, typos, and AS path forging. They claimed a 95.71% detection accuracy on the ground-truth dataset. However, the proposed system cannot characterize tier-1 and tier-2 malicious AS paths or the AS path of a malicious origin with a tier-1 or tier-2 upstream provider. Likewise, hijacking events that involve stealthy sub-prefix methods could easily evade the system. It also misclassified some small valleys (Orange AS5511) as attacking AS.

X. Discussion

A. Validating intentional hijacking.

Although we have validated that the original study's findings mostly still hold, it is challenging to ascertain the intent of hijacks among the alleged serial hijackers. As an example, we studied the behavior of the alleged serial hijacker AS19529 [9]. 41.6% of the prefixes announced by this AS has a 100% match of MOAS with AS20473, AS397460, and AS39879, which technically indicates hijacking events. Further investigation shows a customer-provider relationship between AS19529 and the above ASes, which is technically possible as claimed in [26] but is not a recommended practice. In contrast, AS134190 originates prefixes (45.9%) that belong to about 15 ASes in different regions (ARIN, APNIC, and AFRINIC). A further investigation shows no customer-provider relationships amongst AS134190 and any of the 15 ASes. Although these announcements seem suspicious, we cannot conclude that

these BGP activities are malicious hijacking events. These findings confirm yet again that identifying hijacks is challenging, especially when trying to infer malicious intent.

B. Responding to hijacking

On a different note, even if we could identify serial hijackers reliably, the research and operator community still needs to decide how to deal with them. For instance, Section VIII shows that alleged serial hijackers often remain active or at least allocated by the RIR.

One way forward could be a reputation system to classify ASes based on their prefix origination behavior and make their reputation score public. Not peering with such ASes with a low reputation could force them to follow the existing BCPs, thereby minimizing many of the inconsistencies we have observed and improving the accuracy and reliability of our detection methods. Similarly to [31], the approach of "naming and shaming" could have some positive impact on the security posture of the interdomain routing ecosystem, thereby promoting transparency and accountability.

This approach could enhance the security posture and resilience of the Interdomain routing ecosystem. However, there could be potential legal issues, such as wrongly flagging benign networks (false positives), that could damage the reputation of a network, thereby negatively affecting its chances of competing in the open market. Therefore, we recommend a collaborative effort among academic researchers, the industry, legal practitioners, and all stakeholders involved in the Internet's governance to develop standard frameworks, evaluation metrics, and testbeds to enable research products to be rigorously evaluated and approved before deployment. These stakeholders must also provide a transparent and unambiguous appeal process that quickly resolves false positives. These fair processes would significantly safeguard the reputation of wrongly flagged networks, ensure fair competition in an open market, and enhance the resilience and security of the Internet.

C. Immediate actionable recommendation

Similar to the MANRS initiative [32], we also believed an increased adoption of RPKI's Route Origin Authorization (ROA) and enforcing Route Origin Validation (ROV) would significantly mitigate most hijacking activities (origin hijacks) we observed among the alleged serial hijackers. Therefore, we recommend a collaborative effort among network operators to register ROAs for all their prefixes, including the unused ones, and enforce route validation to enable a more resilient and secure interdomain routing ecosystem.

Similarly, the authors of [26] claimed that it is technically possible for providers to announce their customers' prefixes for some technical reasons, and we have observed such behavior in Section V. However, this practice would technically lead to an origin hijack if the provider did not have ROAs for the announced prefixes. Therefore, prefix owners should always add ROAs specifying that their provider can announce their prefix(es) on their behalf, which would mitigate such false positives.

D. Trade-off analysis and limitations

Despite the efforts to mitigate various caveats, our study still has limitations.

a) Short lived hijacks: Our single daily snapshot sampling method is prone to missing some malicious events, such as short-lived hijacks. Like the original paper [9], our method can only detect prefix origin hijacks. Malicious actors carrying out path hijacks will remain undetected. The trade-off of choosing fewer snapshots and missing short-lived hijacks remains an open challenge for the research community that has to handle the continuously growing amount of routing data.

For example, our exploratory analysis shows that our dataset contains significant redundancy, which could distort our analysis [21]. Therefore, an adaptive sampling method could help us choose a subset of the data to capture crucial events to streamline our analysis and develop efficient detection methods, enabling us to conduct such resource-intensive research with minimal resources. Such an adaptive approach could increase the sampling rate in period where many events on the BGP control plane are reported and could lower the sampling rate during less eventful periods.

In our use cases, however, we mainly investigated the long-term behavior of serial hijackers, which mostly lasted longer than a day. For example, a so-called serial hijacker, AS134190, announced 0.61% of its prefixes for less than or exactly a day, and the remaining 99.39% lasted more than a day. Similarly, AS197426 announced only 0.11% of its prefixes for precisely a day or less, and 99.89% were more than a day. These statistics show that our sampling methods should capture the dominant attributes of serial hijackers, thereby minimizing skewing our analysis and understanding of the serial hijacking activities.

b) False positives: Moreover, our classifier recorded false positives, such as flagging DDoS protection services as serial hijackers, possibly due to MOAS features. We intend to use Recurrent Neural Networks (RNN) and word embeddings like Word2Vec to extract crucial features representing semantic meaning. For instance, the authors of [33] extended the concept of Word2Vec as BGP2Vec that uses AS-paths to embed each ASN into a dense vector and apply it on an RNN to classify BGP hijacks. The authors claimed an accuracy of 99.99% and 0.00% false positives, which has the potential to minimize the false positives of our classifier.

c) Reproducibility challenges: Lastly, although we found the original paper's methodology and findings to be as claimed, reproducing them was nontrivial because of many missing artifacts that we needed to figure out ourselves. Therefore, publishing papers with supporting artifacts should be encouraged. As a result, we are making all our artifacts available, including our data collection and preprocessing scripts, to improve transparency and complement research efforts.[5]

[5] git@github.com:ebrimajaw/revisiting-serial-BGP-hijackers.git

E. Ethical considerations

Our study has no ethical concerns because we used publicly available RIR and BGP data for all our analyses.

XI. Conclusion

In this work, we revisited the work of Testart et al. to validate their findings and investigate the current dynamics of serial hijackers during the past five years (2019-2023). Our findings have validated the majority of the observations of the original study, such as the distinct patterns of prefix origination among regular and serial hijackers. Moreover, we have shown that it is possible to characterize serial hijackers with much less data and developed a competing classifier that evaluated the current trends of serial hijacking activities. We have observed indications of malicious and legitimate unconventional behavior among our flagged ASes and confirmed that there are still ASes listed to be involved in some malicious behavior. Our findings show that there has been no increase in the evolution of serial hijacking activities during the last five years. At the same time, we have confirmed that most alleged serial hijackers from the original study remain active on the Internet, showing that the routing community has not found a common path to deal with malicious behavior. Finally, we have recomputed the ground truth with metadata and will provide the artifacts to reproduce our work when publishing the paper.

Acknowledgment

The authors would like to thank the anonymous reviewers and our shepherd, Lars Prehn, for their valuable feedback on our paper. We gratefully acknowledge the support of Thomas Krenc for facilitating our access to the historical BGP data (2014-2018) at CAIDA. This work was carried out as part of the Network Security program of the Twente University Centre for Cybersecurity Research under grant number 20003215 (TUCCR). Cristian Hesselman's work was also part of the CATRIN and UPIN projects, both of which received funding from the Dutch Research Council (NWO).

References

[1] B. Al-Musawi, P. Branch, and G. Armitage, "BGP Anomaly Detection Techniques: A Survey," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 1, pp. 377–396, 2017. [Online]. Available: http://ieeexplore.ieee.org/document/7723902/

[2] H. Ballani, P. Francis, and X. Zhang, "A Study of Prefix Hijacking and Interception in the Internet," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 4, pp. 265–276, Aug. 2007.

[3] Z. Zhang, Y. Zhang, Y. C. Hu, and Z. M. Mao, "Practical defenses against BGP prefix hijacking," in *Proceedings of the 2007 ACM CoNEXT conference on - CoNEXT '07*. New York, New York: ACM Press, 2007, p. 1. [Online]. Available: http://portal.acm.org/citation.cfm?doid=1364654.1364658

[4] P. Paganini, "Russian telco rostelecom hijacks traffic for IT giants, including google," 2020-04-06. [Online]. Available: https://securityaffairs.com/101134/security/rostelecom-telco-hijacks-internet-traffic.html

[5] R. Hiran, N. Carlsson, and P. Gill, "Characterizing large-scale routing anomalies: A case study of the china telecom incident," in *Passive and Active Measurement*, M. Roughan and R. Chang, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, vol. 7799, pp. 229–238, series Title: Lecture Notes in Computer Science. [Online]. Available: http://link.springer.com/10.1007/978-3-642-36516-4_23

[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, pp. 2471–2486, Dec. 2018. [Online]. Available: https://ieeexplore.ieee.org/document/8481500/

[7] G. Huston, M. Rossi, and G. Armitage, "Securing BGP — A Literature Survey," *IEEE Communications Surveys & Tutorials*, vol. 13, no. 2, pp. 199–222, 2011, conference Name: IEEE Communications Surveys & Tutorials.

[8] A. Siddiqui, "KlaySwap - another BGP hijack targeting crypto wallets," 2022-02-17. [Online]. Available: https://www.manrs.org/2022/02/klayswap-another-bgp-hijack-targeting-crypto-wallets/

[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 *Proceedings of the Internet Measurement Conference*, ser. IMC '19. New York, NY, USA: Association for Computing Machinery, 2019-10-21, pp. 420–434. [Online]. Available: https://dl.acm.org/doi/10.1145/3355369.3355581

[10] J. Schlamp, R. Holz, Q. Jacquemart, G. Carle, and E. W. Biersack, "HEAP: Reliable Assessment of BGP Hijacking Attacks," *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 6, pp. 1849–1861, Jun. 2016. [Online]. Available: http://ieeexplore.ieee.org/document/7460217/

[11] Q. Jacquemart, G. Urvoy-Keller, and E. Biersack, "A Longitudinal Study of BGP MOAS Prefixes," in *Traffic Monitoring and Analysis*, D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, A. Kobsa, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, D. Terzopoulos, D. Tygar, G. Weikum, A. Dainotti, A. Mahanti, and S. Uhlig, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, vol. 8406, pp. 127–138, series Title: Lecture Notes in Computer Science. [Online]. Available: http://link.springer.com/10.1007/978-3-642-54999-1_11

[12] W. Li, Z. Lin, M. I. Ashiq, E. Aben, R. Fontugne, A. Phokeer, and T. Chung, "RoVista: Measuring and analyzing the route origin validation (ROV) in RPKI," in *Proceedings of the 2023 ACM on Internet Measurement Conference*, ser. IMC '23. New York, NY, USA: Association for Computing Machinery, 2023-10-24, pp. 73–88. [Online]. Available: https://dl.acm.org/doi/10.1145/3618257.3624806

[13] A. Azimov, E. Uskov, R. Bush, K. Patel, J. Snijders, and R. Housley, "A profile for autonomous system provider authorization," Internet Engineering Task Force, Internet Draft draft-ietf-sidrops-aspa-profile-06, 2021-07-30, num Pages: 9. [Online]. Available: https://datatracker.ietf.org/doc/draft-ietf-sidrops-aspa-profile-06

[14] Q. Li, J. Liu, Y.-C. Hu, M. Xu, and J. Wu, "BGP with BGPsec: Attacks and Countermeasures," *IEEE Network*, vol. 33, no. 4, pp. 194–200, Jul. 2019. [Online]. Available: https://ieeexplore.ieee.org/document/8594708/

[15] MANRS. (2024) Network operator participants. [Online]. Available: https://manrs.org/netops/participants/

[16] CAIDA. (2023) BGPView. Original-date: 2017-10-05T23:43:21Z. [Online]. Available: https://github.com/CAIDA/bgpview

[17] Number Resource Organization (NRO) , "RIR Statistics," https://www.nro.net/about/rirs/statistics/, 2023.

[18] Spamhaus. (2024) DROP - don't route or peer lists - the spamhaus project. [Online]. Available: https://www.spamhaus.org/drop/

[19] UCEPROTECT. (2024) Blacklist policy level 2. UCEPROTECT®-Network. [Online]. Available: https://www.uceprotect.net/en/index.php?m=3&s=4

[20] Emile Aben. (2024) Route collection at the RIPE NCC - where are we and where should we go? [Online]. Available: https://labs.ripe.net/author/emileaben/route-collection-at-the-ripe-ncc-where-are-we-and-where-should-we-go/

[21] T. Alfroy, T. Holterbach, T. Krenc, K. Claffy, and C. Pelsser, "Internet science moonshot: Expanding BGP data horizons," in *Proceedings of the 22nd ACM Workshop on Hot Topics in Networks*. ACM, 2023-11-28, pp. 102–108. [Online]. Available: https://dl.acm.org/doi/10.1145/3626111.3628202

[22] RIPE NCC. (2016) Routing information service (RIS). [Online]. Available: https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/

[23] University of Oregon. (2023) Routeviews project. [Online]. Available: https://www.routeviews.org/routeviews/

[24] Dyn Blog. (2024) Iran leaks censorship via BGP hijacks. [Online]. Available: https://web.archive.org/web/20170112063331/http://dyn.com/blog/iran-leaks-censorship-via-bgp-hijacks/
[25] CleanTalk. (2024) ASN list - ASN lookup blacklist and IP's spam statistics. [Online]. Available: https://cleantalk.org/blacklists/asn?page=11
[26] K. Z. Sediqi, A. Feldmann, and O. Gasser, "Live long and prosper:analyzing long-lived MOAS prefixes in BGP," 2023-07-17. [Online]. Available: http://arxiv.org/abs/2307.08490
[27] ASRank. (2024) AS rank : GraphQL API overview page. [Online]. Available: https://api.asrank.caida.org/v2/docs
[28] RIPE NCC, "Routing Information Service (RIS)," https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris, 2023.
[29] Y. Xiang, Z. Wang, X. Yin, and J. Wu, "Argus: An accurate and agile system to detecting IP prefix hijacking," in *2011 19th IEEE International Conference on Network Protocols*. Vancouver, BC, Canada: 2011 19th IEEE International Conference on Network Protocols, 2011-10, pp. 43–48, ISSN: 1092-1648. [Online]. Available: https://ieeexplore.ieee.org/document/6089080
[30] S. Cho, R. Fontugne, K. Cho, A. Dainotti, and P. Gill, "BGP hijacking classification," in *2019 Network Traffic Measurement and Analysis Conference (TMA)*. Paris, France: IEEE, Jun. 2019, pp. 25–32. [Online]. Available: https://ieeexplore.ieee.org/document/8784511/
[31] M. Luckie, R. Beverly, R. Koga, K. Keys, J. A. Kroll, and K. Claffy, "Network hygiene, incentives, and regulation: deployment of source address validation in the internet," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019, pp. 465–480.
[32] Network operator actions. [Online]. Available: https://manrs.org/netops/network-operator-actions/
[33] T. Shapira and Y. Shavitt, "A deep learning approach for IP hijack detection based on ASN embedding," in *Proceedings of the Workshop on Network Meets AI & ML*, ser. NetAI '20. Association for Computing Machinery, 2020-08-10, pp. 35–41. [Online]. Available: https://dl.acm.org/doi/10.1145/3405671.3405814