SIDN Labs

https://sidnlabs.nl

Author Version

Title: Reproducibility Study and Assessment of the Evolution of Serial BGP Hijacking Events

Authors: E. Jaw, M. Müller, C. Hesselman, L. Nieuwenhuis

Published in: IEEE Transactions on Network and Service Management (Volume: 23)

Link to original: https://doi.org/10.1109/TNSM.2026.3671613

©2026 IEEE

Reproducibility Study and Assessment of the Evolution of Serial BGP Hijacking Events

Ebrima Jaw, *Student Member, IEEE*, Moritz Müller, Cristian Hesselman* and Lambert Nieuwenhuis

***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 intentional or unintentional illegitimate announcements of network resources that can compromise the confidentiality, integrity, and availability of communication systems. In the past, the 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" has explicitly focused on these repeat offenders, and it 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 reproduced the 2019 study and showed that we can identify potential offenders with less data while achieving similar accuracy. Our study confirms that there has been no significant increase in the evolution of serial hijacking activities in the last five years. We then extend their research, further analyze the characteristics of the serial hijackers, and show that most of the alleged serial hijackers are still active on the Internet. We also find that 22.9% of the hijacks violated RPKI objects but were still widely propagated, and that even MANRS participants were among the propagating networks.**

***Index Terms*—Border Gateway Protocol (BGP), Prefix hijacks, RPKI, Regional Internet Registries (RIR), Serial hijackers.**

I. Introduction

OUR society has increasingly relied on the Internet infrastructure for its day-to-day critical needs over the past decade [1], which consists of a decentralized set of heterogeneous Autonomous Systems (ASes) [2], [3]. An AS is a network under the control of an independent administrative domain that uses the Border Gateway Protocol (BGP) to establish inter-domain routes and exchange reachability information, such as IP prefixes (203.0.113.0/24) and path attributes (e.g., origin and next-hop) with its peers on the Internet [4].

BGP is the widely adopted inter-domain routing protocol that has significantly enhanced the Internet's resilience and scalability. For example, it enables ASes to form a highly decentralized, interconnected, and robust network that efficiently and dynamically routes traffic across available paths, minimizing network failures and ensuring uninterrupted services [4].

Although the decentralized and trust-based nature of BGP has facilitated the scalability and linear growth of the Internet, these attributes also make it vulnerable to misconfigurations and prefix hijacks [5]–[11], often exploited by malicious actors [8], [9], [12]. During such a hijack, an AS illegitimately advertises prefixes of another AS, which is made possible by BGP's lack of default authentication and verification of routes.

*This work was performed while Cristian was affiliated with SIDN Labs.

The authors of [13] classify prefix hijacking incidents into forged AS paths and origin changes. AS path forging is a malicious technique of inserting an attacking Autonomous System Number (ASN) into a regular path originating from a legitimate network (AS) [14], [4], [6], [9]. In contrast, origin changes happen when an AS originates a prefix or a more specific prefix that belongs to another AS, causing apparent changes to the Internet's topology [15]. There have been several reports of prefix hijacking events that attempted to gather information [16], censorship [17], and financial gains [6], [18].

Without countermeasures in place, bogus routes can propagate widely across the Internet. For example, 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 [19]. Similarly, the recent hijacking and route leak incident involving a more specific prefix (1.1.1.1/32) of Cloudflare's 1.1.1.0/24 made the DNS resolver unreachable for over 300 networks in 70 countries for about eight hours before fully resolved [20].

Internet routing security solutions like Resource Public Key Infrastructure (RPKI) have seen a growing adoption. It secures global Internet routing by providing a framework to validate the authenticity of BGP announcements [21], [22]. However, operational challenges such as the need for continuous maintenance of Route Origin Authorizations (ROA) objects, dependencies on external systems, and coordinated efforts among Internet operators to ensure RPKI's robustness and resilience have hindered its effectiveness. [23]. RPKI's Route Origin Validations (ROV) allows ASes to verify the authenticity of the origin of prefixes [24], enabling them to detect and mitigate prefix origin hijacks. Nonetheless, the fact that only a minority of ASes (12% increase during 2022-2023) are enforcing ROV can compromise the security goals of RPKI [25]. As a result of these challenges, the last decade has seen researchers from both academia and industry make significant progress in providing reactive and proactive routing security solutions to minimize the security threats against the inter-domain routing ecosystem [8], [9], [13], [14], [26], [27].

Irrespective of these efforts, prefix hijacks remain a threat to the Internet [18]. Testart et al. have shown in 2019 that certain networks (*so-called serial hijackers*) occasionally perform multiple hijacks repeatedly and for several days, weeks, or even months [8]. Their study found around 900 of such potential *serial hijackers* active between 2014 and 2018.

Since then, serial hijackers have not received any attention from the research community anymore and 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, which is an extended version of a conference paper [28], 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.
- We develop our competitive classifier which relies on fewer features.
- We show that we can characterize and identify serial hijackers with 99.65% less data, demonstrating that studying (serial) hijackers is also feasible for researchers and operators with fewer resources.
- We showed that 22.9% of our flagged ASes' announcements were RPKI-invalid, which many route collector peers see, indicating insufficient enforcement of ROV among many peer networks, including MANRS participants.
- 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.
- Finally, we provided a comprehensive literature review and categorization of the most relevant works on BGP hijacking incidents.

The remaining paper is structured as follows: Section II extends the background information presented in [28]. Section III presents an overview of serial hijackers while IV highlights the details of the used datasets. Sections V to VIII present an updated methodology and a new analysis of our findings. Section X presents an extended discussion and categorization of related work. Finally, in Section XI we discuss our findings, and XII concludes our study.

II. Background

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

A. Core Functionalities

BGP is the essential inter-domain routing protocol that allows network administrators to define technical, business, and regulatory routing policies to handle the Internet's dynamic nature and continued growth [29]. Routers at the edge of ASes use *update messages* to exchange reachability information about IP address space with other ASes' routers, enabling effective, optimal Internet traffic routing [30]. Figure 1 illustrates the core functionalities of BGP among different ASes and the associated security vulnerabilities like prefix origin hijacks. For instance, AS64496 announces 192.168.0.0/16 to AS64499, and we assumed AS64496 hosts critical services on 192.168.0.1. Subsequently, AS64499 appends itself to the AS path and propagates it to AS65536, which updates its routing table to enable end users to reach the services hosted on 192.168.0.1.

The routers of neighboring ASes calculate the best route to reach an IP prefix based on numerous BGP selection criteria, such as path length, preferences, and policies [5]. However,

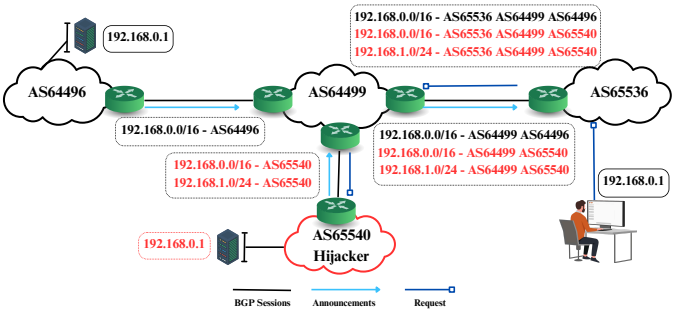


Fig. 1. BGP operations and security vulnerabilities

routers generally prefer routes that claim a path towards a more specific part of a prefix (e.g., preferring a path towards 192.168.1.0/24 over one towards 192.168.0.0/16). Finally, peers prefer routes that provide connectivity via a customer or peer rather than via a transit provider, or that provide a shorter path to a prefix (fewer hops in the routing).

B. BGP Hijacks

BGP lacks built-in authentication and validation mechanisms for connected ASes to verify the authenticity of available routes, causing wide-scale malicious or unintentional hijacks [4], [8], [9], [20]. The two main classifications of BGP hijacks are AS-path and prefix-origin hijacks. In a path hijack, malicious actors can manipulate the path by claiming that the attacker's AS lies on the path toward the victim's AS. If advertised as more specific, this route will likely impact a wide range of ASes.

In contrast, prefix origin hijacks are an illicit origination of IP prefixes belonging to another AS(es) [31]. For instance, Figure 1 shows a prefix origin hijacking event, where AS65540 advertises 192.168.0.0/16 and 192.168.1.0/24 of AS64496, which AS64499 and AS65536 accepted due to the lack of implementation and/or enforcement of filtering. Consequently, these bogus routes tricked AS64499 into routing traffic requests from AS65536, initially intended for AS64496, to AS65540 instead. It is essential to note that AS65540's announcement of 192.168.0.0/16 would lead to a Multiple Origin AS (MOAS) announcement (prefixes originating from both the victim and the malicious network [32]), which might not be very effective depending on many factors, such as AS path lengths and implemented routing policies. However, AS65540's more specific announcements of 192.168.1.0/24 would attract all the traffic meant for AS64496's 192.168.0.0/16 towards itself. Consequently, AS65540 could easily intercept, manipulate, or blackhole AS64496's traffic [6], [18].

C. Countermeasures

BGP hijacking mitigation measures such as Autonomous System Provider Authorization (ASPA) [33] and BGPSec [27] are either in early deployment or have yet to be deployed in the wild. Basic but already deployed countermeasures are RPKI-based Route origin authorization and validation to mitigate prefix origin hijacks. RPKI utilizes cryptographic certificates and digital signatures to enable BGP speakers to validate

the legitimacy and integrity of route announcements among themselves [34]. It enables ASes to create cryptographic objects (ROAs) that bind their resources (prefixes and authorized ASNs) to a public key signed by a Certificate Authority (CA) to authorize the origination of IP prefixes [34]. ROAs also include a maximum-length value, defining the maximum length of prefixes an AS is authorized to announce.

The five Regional Internet Registries (RIRs) and their delegated Local Internet Registries serve as publication points for RPKI objects in repositories that contain ROAs with signed certificates [34]. ASes use ROV to verify that announcements received from their neighbors are authentic and conform to published RPKI ROAs. This validation process ensures ASes accept only valid announcements and drop all invalid announcements to mitigate prefix-origin hijacks.

Finally, even though RPKI was already available during the previous study in 2018, the past few years have witnessed considerable growth in its adoption, partially helping prevent hijacks. For example, RPKI-ROV has increased recently (from 6% to 12% between 2022 and 2023) [25]. However, it only protects against origin hijacks, and the question remains to what extent it has contributed to reducing hijacks in general and serial hijacks specifically. We discuss our findings on how many upstream providers propagate RPKI invalid announcements in Section VIII-D.

III. Serial Hijackers

Testart et al. [8] 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 persistent offenders, we have found reports on the NANOG mailing list (2014-2018) that provide anecdotal support for their claims.[1]

A. Profiling Serial Hijackers

Testart et al. selected a list of ASes from the MANRS initiative [35] 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.

To characterize and subsequently identify serial hijackers, they relied on 3 datasets: *First*, they used CAIDA's BGPview tool [36] to create a longitudinal dataset of mappings between IP prefixes and announcing AS numbers. 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 delegated by which RIR [37]. 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 machine learning (ML) classifier relying on ExtraTrees. This classifier flagged about 900 ASes as potential serial hijackers. Finally, they used other supplementary data to cross-validate the flagged networks, such as SPAMHAUS ASN-DROP List [38] and UCE-PROTECT Level 2 spam blacklist [39]. We refer the reader for a detailed description of their paper [8].

B. Reproducibility: Our Motivation

To our knowledge, Testart et al. are the only ones who studied the behavior of serial hijackers [8], which dates back to 2019. Our goal is to study the development of serial hijackers since then. However, when we tried to reproduce their methodology, we found that the authors relied on terabytes of BGP data, which required storage and processing capacities we did not have.

The increase of data collected at route collectors does not only pose a challenge for us [40]. 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 [41]. 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 would also like to enable other researchers and operators with similar constraints to our own to study this relevant topic. In the end, we want to answer the following research questions:

1) To what extent can we reproduce the findings of the original work using less data?
2) How did serial hijacking events evolve since 2019?
3) How does the deployment of ROA and ROV relate to serial hijacking?
4) What actions did the Internet community take in response to the reported alleged serial hijackers?

Aside from Testart et al. several studies have proposed methods to detect standalone hijacks, andSection X discusses the most relevant attempts.

IV. Datasets

A. Longitudinal BGP dataset

We first demonstrate the redundancy in BGP data using one month of data (December 2021) to support our rationale for selecting a single snapshot for our study. For each day, we processed all updates across all RouteViews collectors. We computed the distinct prefixes and prefix-origin pairs from each day's first snapshot to provide a baseline for our comparison. We then process the remaining snapshots of each day to identify prefix–origin pairs that appear for the first time (*new prefix-origin pairs*) and those that reappear (*redundant prefix-origin pairs*) despite already being present in the first snapshot of the respective day. Figure 2 summarises our results across all RouteViews collectors.

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

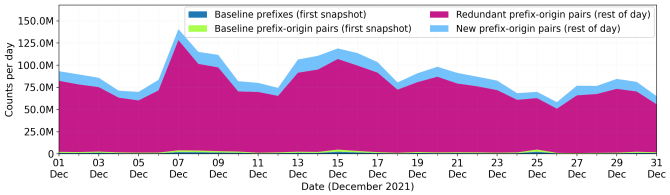


Fig. 2. Redundancy in daily BGP updates across all RouteViews collectors for December 2021. Across all days of the month, we find that the first snapshot of each day already contains 84–90% of all distinct prefix–origin pairs observed that day, while only the remaining 10–16% appear later in the day.

We found that processing snapshots throughout the day yields more distinct prefix–origin pairs than the initial snapshot of each day. However, this increase is primarily due to short-lived or non-persistent announcements that appear once and are never repeated across other days, which could be due to transient traffic engineering [42], as well as more specific announcements [43], [44]. The marginal gain of processing all snapshots in a day is small. For example, only 3–5% of all distinct prefix–origin pairs observed during the whole day are persistent pairs that do not appear in the first 15-minute snapshot. Notably, we observed that 84–90% of all prefix–origin pairs observed during a full day of the month were already present in the first snapshot, yielding a daily average of 86% redundancy in prefix–origin pair records across snapshots of each day. These findings indicate that most of the additional updates observed later in the day are redundant re-announcements rather than new routing information, limiting the marginal benefit of full-day processing. Crucially, this marginal benefit effect arises from how routing behavior is summarized in our analysis. The additional information introduced by later snapshots rarely persists long enough to accumulate into stable behavioral signals. Also, our classifier relies on features aggregated over extended time spans to emphasize repeated, sustained routing behavior at the AS level, such as visibility and stability. In contrast, non-persistent or short-lived announcements typically occur in isolation and do not recur across days or months, contributing negligibly to these aggregates. As a result, including all 288 daily snapshots primarily increases the number of observations without changing the feature distributions that distinguish serial hijackers from legitimate ASes.

Additionally, Jaw et al. [45] recently showed that most BGP announcements in daily BGP traces are dominated by high-frequency, repeated announcements with unchanged attributes that contain little or no new routing information. For instance, they showed that in some BGP update files, fewer than 2% of prefixes account for over 90% of all updates, and a single session contributed an average of 28.67% of the updates over an entire month, mostly by re-announcing the same information. These findings are consistent with our analysis of the entire month of BGP traces (Figure 2), indicating that a significant share of the five-minute snapshots in a day mostly re-announce exact prefix-origin pairs, with limited additional information for distinguishing potential serial hijackers from normal networks.

Based on these findings, we are confident that selecting a single five-minute snapshot per day is sufficient to characterize the attributes of both regular and malicious networks, as shown in Section V.

Since the peer-pfx-origins dataset is only available until the end of 2018, we use CAIDA's BGPView tool [36] to compute our current five-minute snapshots of the peer-pfx-origins dataset for 2019-2023 using BGP updates from RIPE RIS [46] and Routeviews [47].

B. MOAS dataset

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

C. RIR delegation files

Delegation files are publicly available [37]. These files document the allocation of Internet resources, including IPv4, IPv6, and ASNs. We downloaded a single snapshot of these files at the start and end of each year over the five years for each RIR to calculate 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 between the first and last snapshots of the year, due to new records appearing in the latter. Despite these differences, we argue 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 and RIR dataset, which will be available upon request due to their large file sizes.

D. AS paths dataset

Lastly, we downloaded several RIB snapshots from the top five major collectors of Routeviews and RIPE-RIS between 1/1/2019 and 31/12/2023 to provide a comprehensive view of the routing table. We collected the RIB dumps three times per month, on the 1st, mid, and last day of each month, across five years, and parsed the data to extract their timestamps, prefixes, AS paths, and origin ASes as our AS path dataset. We aggregated our data from all collectors into a single daily snapshot to mitigate redundancy caused by potential duplicate views across collectors, and only selected the AS paths that contain our invalid announcements for our analysis. We have removed all incomplete or corrupted files from the downloaded RIB dumps during the aggregation step. However, the ratio of corrupted RIB dumps was negligible, and excluding them

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	

would reduce visibility of our prefixes by only a tiny fraction from the affected collector's viewpoint. We still observed AS paths propagating our prefixes of interest from a significant share of the selected collectors. For example, we observed valid RIB dumps from at least 99.92% of the collectors we selected during our aggregation phase, indicating that the corrupted files constitute a negligible share of our data and that their exclusion would not affect the overall AS paths dataset or our final results. Subsequently, we ended up with 180 snapshot files across the five years. Table II shows the five-year aggregated statistics of AS paths, unique ASNs, and IPv4 and IPv6 prefixes.

TABLE II
DATASET PROPERTIES: AGGREGATES OF THE LONGITUDINAL AS PATHS.

Details of AS Paths Dataset		
Start Date: 01/01/2019	End Date: 31/12/2023	
	IPv4	IPv6
Unique AS paths	198,364,926	43,951,429
Unique prefixes	2,105,659	454,933
Unique ASNs	54,239	15,965

E. Processing

We follow the same data-preprocessing steps described in [8], 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 time of day by comparing results from different single-snapshot timestamps within a day. Each snapshot lists the observed prefixes, 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 [8]. 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.

a) Ground-truth data: Finally, we used the same ground-truth dataset from the original study for consistency and validity, comprising 23 serial hijackers and 217 legitimate ASes. We also adopted the metric from the original study to select 10,623 ASes that announced at least 10 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

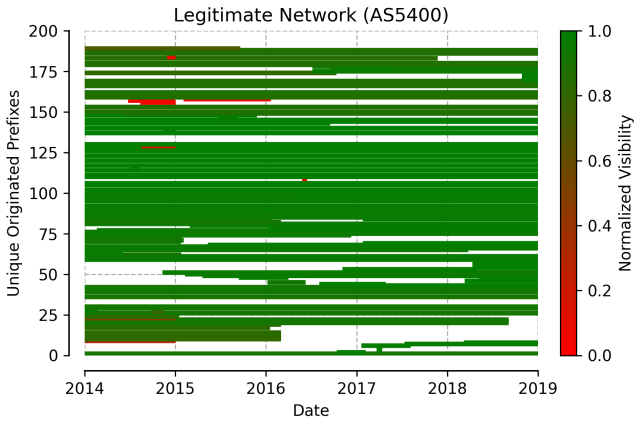
To test our hypothesis that using less data can capture the characteristics of serial hijackers, we compare the behavior described by Testart et al. with our observations. We find that the most highlighted attributes in the original paper hold while using a significantly sparser dataset.

1) Prefix origination patterns: We show in Figure 3 which prefixes the ASes have announced over 5 years and how many peers have observed those announcements, using only one five-minute daily snapshot. Each row represents one prefix. Figure 3a shows the colored normalized visibility of prefix announcement patterns of a legitimate AS, while Figure 3b shows the pattern of a serial hijacker [8].

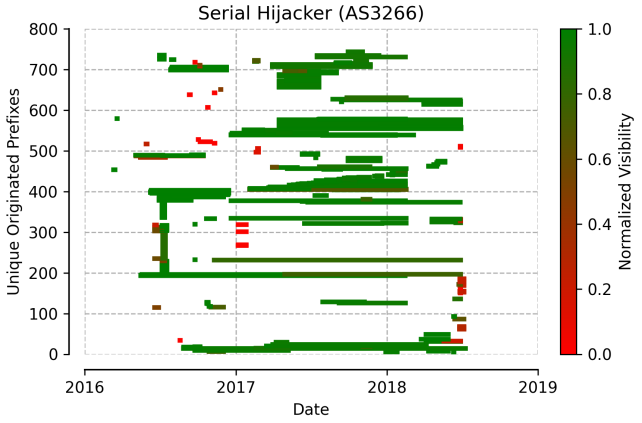
In Figure 3a (and as shown by the original paper), we can see that British Telecom (AS5400) shows a relatively stable origination behavior during the five years, such as consistent prefix announcement duration (33.68% of active prefixes for the entire five year period) and moderately few prefixes compared 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 [8].

In contrast, AS3266 shown in Figure 3b exhibits inconsistent origination patterns over three years, with most prefixes announced for very short periods. For instance, only 52.35% of the prefixes were active for only 1 year (2017-2018), and 0% were continuously active for the entire 5-year period. Moreover, AS3266 showed erratic behavior in the global routing table, with about 800 prefix announcements over only three years, about 75% more than AS5400 over five years.

As with other hijacks, it is hard to say with certainty whether a serial hijacker is acting maliciously or whether misconfigurations are causing the hijacks [48], [49]. 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.

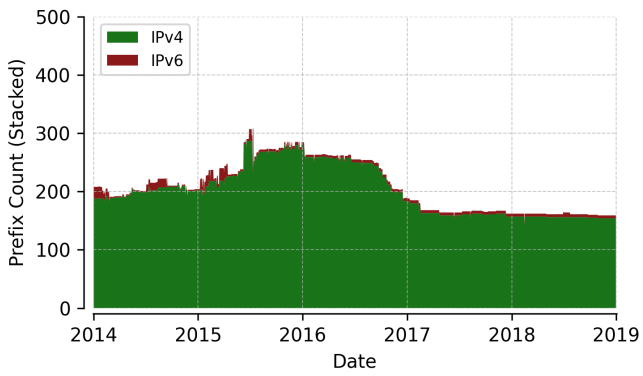


(a) Prefix origination pattern of British Telecom (AS5400).

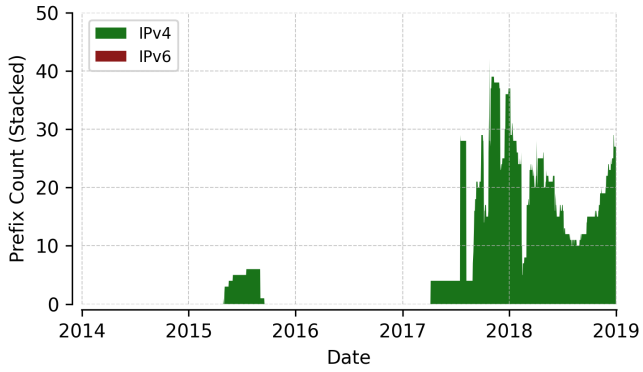


(b) Prefix origination pattern of a serial hijacker (AS3266).

Fig. 3. Reproduced five and three years prefix origination patterns of a legitimate AS (5400) and serial hijacking AS (197426) [8]



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



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

Fig. 4. Reproduced prefix origination changes for a normal (AS7922) and malicious (AS133955) networks over the period of 5 years. [8]

2) Prefix origination changes: Figure 4a shows a consistent origination behavior of AS7922 over five years. Our study missed the spike in Figure 4a in the original study. We confirmed this was due to the deaggregation of larger prefixes and was not related to hijacking.

In contrast, Figure 4b 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 5a 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 [8]. In contrast, Figure 5b 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 AS7922 prefixes were announced for fewer than 50 days. This behavior happens due to local engineering practices we have investigated and validated as explained in (Section V-A2).

VI. Detecting Serial Hijackers with Sparse Data

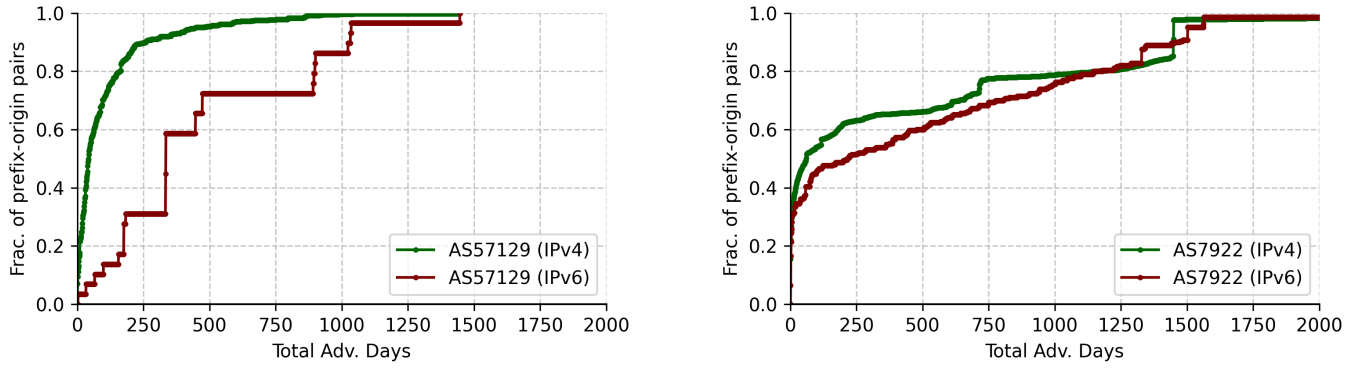
After showing that we can characterize serial hijackers with less data, we now want to *automatically detect* serial hijackers.

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 complete data set. To do so, we use 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 7 illustrates the original classifier's Mean Out-of-Bag (OOB) error on our sampled data. As in the previous study, we use different sampling methods to address the unbalanced training set. We observed that increasing forest size reduces the mean OOB error across all sampling methods (and thereby improves OOB accuracy). Finally, our correlation analysis shows that some features were redundant and did not add

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



(a) Serial hijacker (AS57129)

(b) Normal network (AS7922)

Fig. 5. 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).

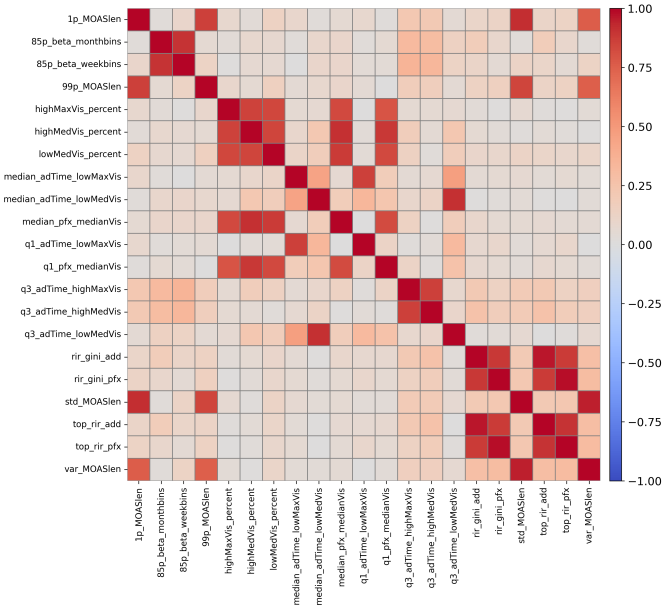


Fig. 6. Correlation dynamics of the subsets of highly redundant features. We observed a near-duplicate structure across many of the features of percentile-based visibility metrics and MOAS length statistics, indicating substantial redundancy in the original features.

much value to 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.

B. Rationale for Feature Subset Selection

We computed the Pearson correlation matrix across all 52 engineered features from the original study to identify potential redundancies, thereby reducing model training time [50] and enhancing model explainability and performance [50], [51].

Our correlation analysis shows that some features were redundant and did not add much value to the final prediction. For readability, Figure 6 shows the absolute correlation between the subsets of features that are highly correlated (*features with* $|\rho| \geq 0.85$) with at least one other feature. Our results reveal several blocks of strongly correlated features. For instance, we observed strong correlations among features derived from quantile statistics of advertisement time, as expected, given that these features are different summaries of the same underlying temporal behaviour of announcements from the involved ASes.

Similarly, the MOAS-length features all summarize the same underlying set of MOAS lengths for the involved ASes. For instance, if the MOAS lengths of an AS become more variable, larger, or more concentrated, all computed statistics, such as minimum, maximum, and percentiles, shift in predictable ways because they are derived from the same values; hence, the strong correlation of the MOAS-based features. Additionally, the prefix-visibility metrics that quantify how widely an AS's prefixes propagate across vantage points tend to rise and fall together because they capture different summaries of the same underlying visibility distribution. For instance, an AS with highly visible prefixes will simultaneously exhibit a high median visibility, a large fraction of highly visible prefixes, and high visibility percentiles. In contrast, if it has low-visibility prefixes, all these metrics will drop in parallel. This shared dependence on the identical prefix-visibility distributions explains why this group of features shows strong pairwise correlations, thereby leading to redundancy among the prefix-visibility-based features.

Based on these insights into feature redundancy, we selected only 44 of the original 52 features for training and evaluation. Notably, some of the final 44 selected features remain correlated with at least one of the remaining features. However, this does not introduce unnecessary redundancy or overfitting in our experiments for two reasons. First, our adopted ExtraTree classifier is inherently robust to correlated inputs because feature selection occurs implicitly through randomized splits and ensemble averaging [52], [53], preventing any single redundant feature from dominating the model. Second, we have verified that removing the highly correlated features does not reduce our model's performance. For instance, when reducing the number of features from 52 to 44, we remove only correlated features whose exclusion does not affect predictive

performance. It is crucial to note that some of the final 44 selected features are correlated with at least one other selected feature. However, we intentionally retained these correlated features because removing them led to a noticeable drop in model performance. This feature-selection process and our observations indicate that feature correlation alone was not a criterion for removal; instead, we retained features based on their contribution to stable, consistent predictions. Additionally, the consistent precision and recall observed across different sampling strategies and forest sizes in our experiment further indicate that the classifier's performance is robust and not dependent on a specific feature combination and that feature correlation does not lead to overfitting, suggesting that the retained features contribute distinct information relevant to long-term AS behavior.

C. Improving on the original classifier

Even though the original classifier showed good performance on both complete 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 whether we can further improve the classifier. 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 7, we also implemented the Proximity Weighted Random Affine Shadow sampling (ProWRAS) and Adaptive Synthetic Sampling (ADASYN) oversampling techniques to address the class imbalances within our ground truth dataset. To ensure that this imbalance handling does not compromise generalizability, we apply oversampling strictly during training and evaluate performance using precision–recall–based metrics that emphasize detection of rare hijacking events rather than majority-class accuracy. Moreover, the selected features characterize long-term, persistent routing behavior rather than short-lived activity or frequency-based signals. This approach prevents the model from learning dataset-specific majority patterns introduced by oversampling. Consequently, class imbalance mitigation improves sensitivity to rare events without altering the underlying decision boundaries, allowing the model to generalize consistently across datasets with different class distributions.

TABLE III
PERFORMANCES OF THE TWO CLASSIFIERS AND CROSS-VALIDATION APPROACH

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

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

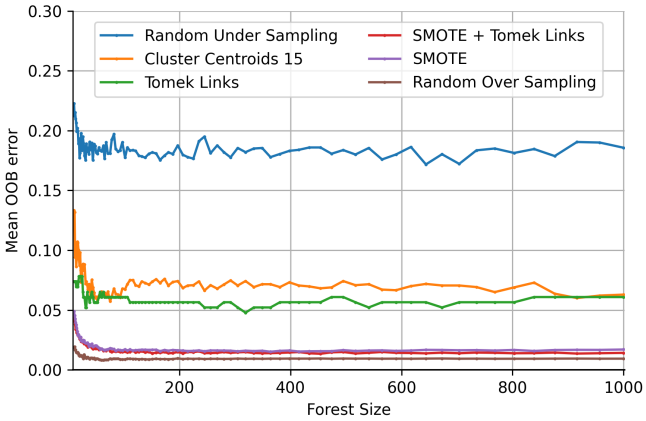


Fig. 7. The mean OOB error of different sampling methods and forest size of the original machine learning classifier on our sparse dataset.

VII. SERIAL HIJACKERS 2019 - 2023

We applied our classifier to 10,622 ASes that originated at least 10 IPv4 prefixes from 2019 to 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 flagged 279 ASes (36.41% of our flagged networks) as exhibiting the characteristics of a serial hijacker, confirming reasonable agreement between the 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. We have computed various metrics presented in Table III to provide a detailed performance assessment of our final model.

B. BGP misconfigurations

A small fraction of the flagged ASes may have resulted from BGP misconfigurations, such as the leaking of private ASNs and fat-finger errors [13] due to path prepending [8]. 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 [17] and reports of single-digit ASes like AS6 and AS2 announcing prefixes belonging to other networks[4]. We excluded the private and

[3]We only reported the precision and recall as done in the original study
[4]https://seclists.org/nanog/2018/Apr/206

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 [8], 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 that our classifier flagged 14 ASes (1.8817% of our flagged networks). The manually compiled list is small, and we could miss other networks. Additionally, these highly dynamic ASes that operate large-scale DDoS mitigation or traffic-scrubbing infrastructures pose a challenge to our approach. Such ASes may exhibit frequent, rapid, and large-volume routing changes that resemble attack-like behavior when viewed solely through control-plane dynamics. In these cases, repeated prefix re-origination, shifting visibility patterns, or transient MOAS events can be driven by legitimate mitigation workflows rather than malicious intent. Our study does not attempt to explicitly distinguish between serial hijacking behavior and operationally induced routing dynamics associated with DDoS defense. Addressing this distinction would require incorporating additional context, such as mitigation service metadata, customer–provider relationships, or traffic-level signals, which fall outside the scope of our study. We therefore treat these cases as open research questions for future work.

D. Listed flagged ASes

We used SPAMHAUS's Don't Route or Peer Drop List [38], Clean Talk-listed ASNs [54], and UCEPROTECT's Level 3 list [39] 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 Level 3 list of ASes and 8 ASes on the SPAMHAUS Don't Route or Peer Drop List. Although the share of our flagged ASes in the above lists is lower than the claims of the original study, we can still observe potential malicious activity among our flagged networks.

E. Sign of spamming among flagged ASes

We used 1,360 prefixes from Spamhaus DROP, DROPv6, and extended Drop (eDROP) over the duration of our study to identify spamming activity among our flagged ASes. These lists contain RIR allocations and sub-allocated blocks of hijacked prefixes or malicious netblocks [38]. We mapped the timestamps of our daily BGP snapshots to the timestamps of Spamhaus's lists. Then we checked whether any of our flagged ASes announced any listed prefixes at 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 distinguish legitimate networks from serial hijackers. Using MOAS features to classify malicious announcement activities is a logical decision. However, while creating our MOAS dataset, we observed a significant increase in its legitimate use. 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% more often than the ASes 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 for announcing MOAS to improve the classifier [55].

G. Fate of the remaining flagged ASes

We used CAIDA's ASRank API [56] to determine the status of our remaining flagged ASes. Our findings show that only 4 ASes (AS49121, AS264009, and AS263481) have at least one peer among the remaining 96 flagged ASes, while 91.94% of the ASes have no peering relationships. About 50% of these ASes currently have no provider. Our findings confirmed that they disappeared 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. Comparative Analysis: Flagged and Non-flagged ASes

Earlier, we found that a significant number of hijacked prefixes were RPKI-invalid. This raises two questions: First, why would a potentially malicious hijacker announce RPKI-invalid routes at all, given that ROV deployment can substantially reduce their propagation and impact? Second, why did the invalid announcements still propagate widely?

A. Approach

First, we want to understand if flagged ASes announce invalids more often.

Because the non-flagged ASes largely outnumber the flagged ASes, we only select a subset of non-flagged ASes for our analysis to achieve a fair comparison. Our subset of non-flagged ASes announces a similar number of prefixes as the flagged ASes.

We compute the unique prefix counts for all flagged and non-flagged ASes to select a subset of non-flagged ASes for our RPKI-invalid comparison. For each flagged AS, we select an equal number of non-flagged ASes that announce the same number of prefixes compared to our flagged ASes. If we cannot find an exact match, we choose a candidate AS with the smallest absolute difference within a tolerance of 2 prefixes (either greater or less). Once we have a match, we exclude the selected AS from our candidate list to avoid duplicates. Then we used RPKI snapshots from 1/1/2019 to 11/31/2023 to assess the RPKI status of their announcements [57], [58].

B. Announcing invalids

The selected non-flagged ASes only had 385 (0.35%)[5] invalid announcements, which is 97.97% less than the flagged ASes or 49 times fewer invalid announcements compared to flagged ASes (about 19K RPKI-invalids) [58]. We also observed that even though the flagged ASes have a comparable number of unique prefixes, they made 3.75% more announcements than the non-flagged ASes, which aligns with the erratic behavior of serial hijackers discussed in [8].

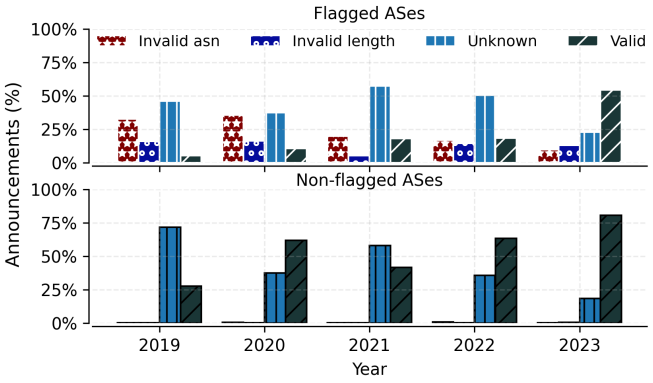


Fig. 8. Aggregates of RPKI status distribution of announcements from flagged and non-flagged ASes

Our results show that the non-flagged ASes' announcements were, on average, valid for five years. Here, we observed a positive trend: in 2019, 27.5% (2.8K) announcements were RPKI valid and 71.8% (7.3K) unknown. In 2023, unknown announcements decreased to 2.2K, and valid announcements rose to 81% (10K). These trends confirm that networks are increasingly creating ROAs for their prefixes. Figure 8 further highlights commitment to RPKI adoption from the non-flagged ASes: They announce significantly fewer invalids than flagged ASes and largely deploy ROAs. This is desirable to enhance the security and stability of the routing infrastructure.

To determine whether the above distinct announcement patterns among our flagged and non-flagged ASes are statistically significantly different, we performed a Chi-square test. We computed our expected values using both networks' announcement ratios and the RPKI status averages. The resulting chi-square statistic ($\chi^2 = 47.92$) exceeded the critical value ($\Upsilon = 7.815$, $df = 3$, $\alpha = 0.05$), leading us to reject the null hypothesis ($p < 0.001$), thereby confirming a statistically significant difference between flagged and non-flagged ASes.

Finally, we acknowledge that RPKI misconfigurations can cause legitimate announcements to appear invalid. However, these cases are usually short-lived and do not persist for months or years or accumulate into the sustained patterns required by our methodology. For example, we assumed that operational-induced ROA misconfigurations, such as incorrect ASNs or maxLength settings, usually persist for hours or days but not for months or years because network operators rapidly correct such issues once routing reachability problems arise. In contrast, our classifier targets ASes that repeatedly and consistently originate suspicious prefixes over extended time windows, which aligns with our observations of the RPKI invalid announcements from our flagged ASes. As a result, while transient ROA misconfigurations may appear in the data, they do not lead to false identification of serial hijackers, ensuring robustness against occasional RPKI errors.

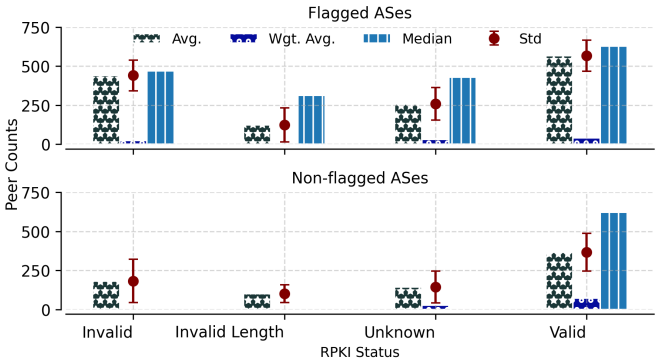


Fig. 9. Descriptive Statistics of peer counts among the RPKI statuses of announcements for flagged and non-flagged ASes.

C. Connectivity

Figure 9 provides more insights into the propagation of the announcements from flagged and non-flagged ASes. We observed that the selected non-flagged ASes predominantly have higher peer counts for their valid announcements than for their invalid announcements. In comparison, flagged ASes have higher weighted-average peer counts of 30.7% and 24.1% for unknown and invalid announcements [58]. Moreover, the low peer counts for invalid announcements among the non-flagged ASes indicate that these announcements were not widely propagated on the Internet compared to valid announcements, possibly because their upstream providers performed route origin validation. Despite moderate averages among the non-flagged ASes, we observed low median counts of their invalid announcements. This indicates that while most peers drop these invalid announcements, which leads to low peer counts, a small set of non-flagged ASes with higher peer counts are propagating them.

In contrast, our results show that, on average, the flagged ASes have higher peer counts for invalid announcements [58]. Nonetheless, we observed lower variance in their invalid

[5]An average of 77 (0.4%) and 54.2 (0.3%) of invalid and invalid length announcements, respectively, during the five years.

announcements than in those of non-flagged ASes, suggesting that flagged ASes exhibit a more uniform distribution of peer counts. This could be due to the limited number of upstream ASes propagating invalid announcements along their AS paths, which we investigated and discussed in the subsequent sections.

D. Propagating invalids upstream

Henceforth, we refer to the list of ASes on the AS path propagating RPKI-invalids from the origin AS as upstream ASes.

We extract all the AS paths that have propagated at least one of the invalid announcements since 2019 from the new dataset discussed in Section IV-D. Subsequently, we evaluated the AS path length and number of upstream ASes for invalid announcements from 2019 to 2023. Then, we used Jackard Distance to measure how frequently these invalid announced prefixes change their upstreams (churn).

For each invalid announcement, we extracted the upstream ASes that propagated it further. In total, we observed about 4.5K upstream ASes (6.37% of the active ASes in the global routing table) that propagate RPKI invalids. About 129 upstream ASes propagated 10% of the invalid announcements. Additionally, we observed that about 204 upstream ASes accounted for 75% of the RPKI invalids (with a median of 183 upstream ASes), indicating that a relatively small number of upstream ASes (204) are responsible for propagating the vast majority of RPKI invalids. Advocating for these few ASes to implement effective security measures, such as IRR or RPKI filtering, would significantly reduce the propagation of these invalid announced prefixes.

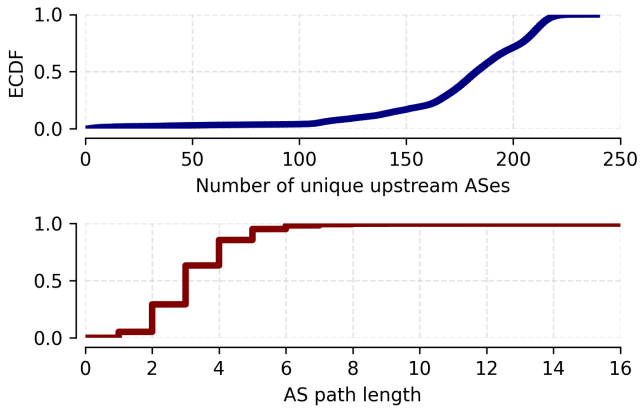


Fig. 10. Distribution of upstream ASes and AS path length for invalid announcements among the flagged networks.

Next, Figure 10 (bottom) shows that about 80% of the invalid announced prefixes have a path length of 4 or less. While shorter AS paths usually indicate normal and efficient BGP routing behavior, malicious networks can exploit these short paths by failing to filter RPKI-invalids to quickly and effectively propagate invalid announcements, leading to a significant negative impact. This threat could exacerbate collateral damage when short AS paths include well-connected ASes, enabling the rapid propagation of these invalid announcements.

In contrast, we observed that about 15% of the RPKI invalid prefixes have a path length of 5-16, which is rare in an ideal BGP operational network [59], [60]. These longer AS paths could be attempts to exploit non-filtering AS paths to evade detection or to drop these invalid announcements; alternatively, they could be caused by BGP rerouting around ASes that drop invalid announcements. Moreover, the authors of [60] claimed that, since the second step in BGP path selection is the shortest AS path, malicious actors could craft longer AS paths to ensure their announcements did not propagate far on the Internet, thereby avoiding public route collectors from reporting their events, which aligns with our observations.

Finally, we used the Jaccard distance to evaluate the daily upstream changes (churn) among the invalid announced prefixes. Here, we test whether flagged ASes switch upstreams more frequently, e.g., because of measures by their former upstreams that limit the impact of their hijacks. A set of upstream ASes for each prefix on a given day is compared with previous sets to identify potential upstream changes. A churn rate near 0.0 means no or minimal upstream changes, and a churn rate near 1.0 signifies frequent upstream changes for given prefixes over time. Our results show that over 90% of invalid announced prefixes are below 0.2.

Specifically, 10% of these prefixes have a churn rate below 0.016, indicating minimal daily upstream changes. However, less than 10% of these invalid prefixes have a churn value near 1.0, signifying that these prefixes almost change upstreams between consecutive days of our measurement period. We acknowledged that such upstream instability could be due to route flapping or misconfigurations. However, potentially malicious networks could exploit such behavior to diversify the propagation of invalid routes or evade detection.

E. MANRS participants' involvement

Mutually Agreed Norms on Routing Security (MANRS) aims to enhance Internet routing security. They encourage participant networks to follow recommended actions, such as *filtering customers' announcements*, to mitigate security incidents, such as route leaks and prefix origin hijacks [61], [62]. This section aims to evaluate the complaint level of 1,356 network participants, extracted from their published lists of participating ISPs and CDNs. Here, we count the number of MANRS participants on an AS path toward invalid announcements. We ensure we count ASes only after they join MANRS, and only after we observe the invalid announcement. Then, we aggregate our observed invalid propagations by date to compute the number of daily invalid announcements, daily fractions of invalid announcements propagated by MANRS participants, and their ratio.

Our results show that of the 1,356 selected MANRS participants, 1,081 ASes (about 80%) were never observed on any AS paths propagating these invalid announcements. This indicates the practical implementation of routing security best practices, such as ROV, among most MANRS participants. However, we found that 262 ASes (19.33%) of our MANRS participants lie on AS paths that propagated our invalid prefixes. That is almost one in every five participants, which

is inconsistent with MANRS's Action I and indicates gaps in adherence to the MANRS policy on filtering customers' invalid announcements. We observed a consistent 18% monthly average fraction of invalid announcements, in which at least one MANRS participant lies on an AS path that propagates these invalids over the five years. On average, 82% of the invalid announcements were propagated by non-MANRS participants.

While MANRS membership signifies a commitment to routing best practices, our observations indicate that MANRS-participating ASes may still propagate routes that appear RPKI-invalid. However, this behavior should not be interpreted as non-compliance or malicious intent, as it may be due to well-known operational factors, including the deployment of partial RPKI validation, delayed synchronization with updated ROA repositories, reliance on third-party transit policies, or the continued acceptance of invalid routes for reachability or traffic-engineering purposes. Additionally, MANRS focuses on encouraging responsible operational practices rather than enforcing strict filtering guarantees. Consequently, the propagation of invalid routes by MANRS members highlights the gap between policy intent and operational reality, rather than a failure of the initiative itself. Therefore, our analysis treats MANRS participation as a valuable indicator of expected benign behavior, while acknowledging that we might still observe RPKI-invalid announcements in practice for many reasons. Understanding when and why such announcements persist remains an open operational question that warrants further investigation.

F. Case studies: Announcing invalids without malicious intent

Our preliminary findings show that business relationships and operational practices, such as customer-to-provider relationships and IP leasing, could be among the potential reasons for RPKI invalids among some flagged ASes. The subsequent sections summarize our two different case studies:

1) Customer-provider relationships: Studies show that providers can announce their customers' prefixes for different operational reasons [55]. Hence, we investigate the announcement behavior of AS19529 (*one of our flagged AS*) to evaluate the impact of such operational practice on RPKI-invalid announcements. Our measurement found that AS19529 originated 9,404 (56.3%) prefixes with a valid ROA record. In contrast, the same AS announced 6956 (41.6%) RPKI invalids. All of these invalids have ROAs for AS20473, AS398795, and AS397460: 6500 (93.36%), 450 (6.48%), and 6 (0.086%) prefixes, respectively. We randomly selected some of these prefixes, manually inspected their path and route objects, and found out all three ASes were customers of AS19529, and some prefixes were distributed using anycast. We speculate that the lack of proper RPKI documentation, due to operational constraints or oversights, could explain some of our legitimate RPKI-invalid announcements [48].

2) IP leasing: The exhaustion of the IPv4 address space has driven rapid growth in the IP leasing industry, enabling cloud providers and content-delivery networks to meet the needs of their growing customer bases [63]. Depending on the leasing model, the origin AS of the leased address space could change frequently. This could lead to RPKI invalids if an ROA object covers these prefixes but is not updated correctly or on time, undermining the effectiveness of current BGP hijacking detection methods that rely on RPKI data.

In this analysis, we used 47.3K inferred leased prefixes (collected by [63]). The goal was to quantify whether our observed RPKI invalids were due to leasing [63]. The list of IPs dates back to April 2024, but our data ends in November 2023. We assume that IPs available for leasing in 2024 were also available for leasing in the second half of 2023. Thus, we only studied the last five months (July to November 2023) of our flagged RPKI invalid announcements. Then, we considered our RPKI invalids as part of the leased prefixes only if their announcement dates were on or after the *last-modified date* field of the leased prefix(es), to minimize counting invalid prefixes before their announcement dates. Note that we count invalid entries that match the leased prefixes when there is no *last-modified* date, which could lead to overestimation. Subsequently, we found 44 of our RPKI invalid announcements among the leased IPs (0.093% of inferred leased IPs and 0.2316% of the invalid announcements).

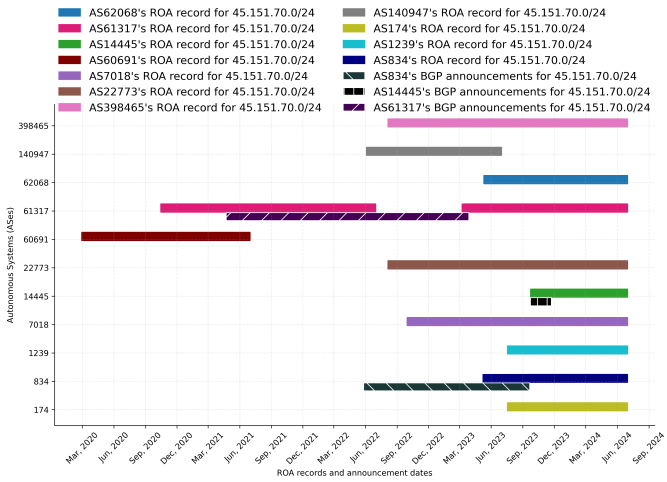


Fig. 11. Occurrences of our RPKI-invalid prefixes observed among the five RIR-leased IP addresses across five years

As an example, Figure 11 shows the timelines of BGP announcements for 45.151.70.0/24 and its corresponding ROA records for all ASes authorized to announce the prefix. Three different ASes announced 45.151.70.0/24 between 2020 and 2024, with 11 different ASes having valid ROA records for it. The announcements by AS834 and AS61317 were RPKI invalid, which is justified in the observed BGP announcements outside their ROA validity records as shown in Figure 11.

AS834 (IPXO LLC), an IP leasing company [64], already announced most of the above 44 RPKI invalid prefixes before having a ROA record. At the same time, conflicting ROA records existed. Using CAIDA's ASRank dataset, we checked customer-provider relationships between AS834 and any ASes holding valid ROA records for 45.151.70.0/24 and found no customer-provider relationships. So, we speculate that these RPKI-invalids could be due to IP leasing and do not have malicious intent.

IX. 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 [65] and archives containing information about IP address space allocations and assignments of the five regional RIRs [37].

Of the 23 ASes, 9 (39.0%) were still active by the end of 2022. Their AS number was continuously allocated, even after they were accused of hijacking Internet resources and were still announcing prefixes. The second-largest group is ASes, which were not allocated for a 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 from the one carrying out the hijacks. 4 ASes (17.4%) are no longer allocated 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 as of now, we find that these ASes are occasionally or continuously announcing MOAS. At some point, one AS announced almost 20 times as many MOAs as 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 for a hijacked AS to be de-allocated varies. ASes remained assigned for a median of 239 days after their last BGP announcement, but one AS remained allocated for 1,609 days. In contrast, two ASes announced prefixes even though their ASs were no longer allocated for 20 and 24 days, respectively. 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 reason could be that the community did not punish them for their malicious behavior.

X. Related Work

The research community has proposed different approaches to detecting and classifying routing hijacks. This section presents a comprehensive extension of the original paper's related works [28].

Section X-A presents research efforts on real-time monitoring, detection, and mitigation of BGP hijacking. Section X-B outlines related works using an RPKI-based approach. Section X-C summarizes investigative methods of AS hijacking incidents for spamming and analyses of the underlying root causes of prefix hijacking incidents. Section X-D presents literature on SubMOAS event classification and analyses of hijacking incidents, and Section X-E summarizes various ML-based methods for prefix hijacking and origin change detection.

A. Prefix hijacking monitoring, detection & mitigation methods

The authors of [26] proposed a Prefix Hijack Alert System (PHAS). Their technique provides reliable, real-time notifications of potential prefix hijacks to the legitimate owners of those prefixes. The authors leverage BGP feeds from collectors to construct origin sets (the number of ASes originating a prefix) for each prefix, detect origin hijacks, and use mailing services to notify the victim swiftly.

Zhang et al. [66] presented iSPY. This approach enables end users to consistently monitor their network reachability from external transit networks. iSpy could detect prefix hijacking events that lead to connectivity loss, but could not detect hijacks of unannounced prefixes.

Xiang et al. [67] 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 about 12K anomalies and detected 63 possible hijacking incidents. However, it can only detect hijacks that led to blackholing, potentially missing non-blackholing serial hijacking events. Argus was not validated to detect other sub-prefix hijacking techniques. Finally, Argus has a delay of $\approx$10 seconds and requires a waiting period before an initial detection of a suspicious hijacking incident.

Finally, Sermpezis et al. [9] 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 targeting other prefixes outside the network and experiences delays accessing the live BGP feed, which could affect performance.

B. RPKI-based prefix hijacking evaluation

Wählisch et al. presented an early study of BGP updates against the RPKI to examine the occurrence of BGP hijacking incidents [49]. The authors attempted to understand the root causes of *valid* and *invalid* prefix announcements using a month's worth of RIPE RIS RRC00 table dumps, which contain 432,864 prefixes. The study concluded that most invalid announcements were due to technical and operational issues related to RPKI adoption, including misconfigurations, except for one event that lasted 37.8 hours. However, their study dates back to 2012, when RPKI covered only 2% of the global routing table prefixes.

C. Investigative method of AS hijacking and the root causes of prefix hijacking events

The authors of [7] investigate malicious AS hijacking for spamming activities. They investigate the events using control-plane and data-plane sources, enabling them to reconstruct the hijacking details. Their approach reveals how the attacker deceitfully used a forged authorization letter to claim ownership of all the prefixes of AS31733, thereby evading detection from its upstream providers. The authors derived guidelines for forensic studies of AS hijacking and presented a conceptual early warning system to mitigate such events, but never implemented it.

Vervier et al. [68] adopted an approach that combines spam traps, netflow data, traceroutes, IRR, and BGP data to determine the root cause of prefix hijacking events. The authors concluded that the hijacker might have legally rented the identified hijacked prefixes for malicious activities such as spamming and suggested using multiple data sources to study hijacks to reduce the high false-positive rates of existing detection methods.

D. Classification of SubMOAS events and analysis of BGP hijacking incidents

The authors of [69] used the RIPE database, a reasoning algorithm, and SSL/TLS scans to propose a filtering-based system that distinguishes malicious and legitimate subMOAS events. The system observed 8,071 subMOAS events and classified 46.5% (legitimate) and 53.5% (potential sub-prefix hijacking incidents). This approach has significantly reduced the search space for sub-prefix hijacking. However, their data source is limited to RIPE, covering only 60% of subMOAS events, which challenges the system's generalizability.

Schlamp et al. [31] extended [69] and proposed a Hijacking Event Analysis Program (HEAP) that utilized numerous datasets, including all five IRR databases, to minimize the false alarm (FA) rates of existing techniques. The authors first derived organizational and business relationships among the parties of the events. They then rule out legitimate events using topology-based reasoning and use internet scanning to identify non-malicious events. Subsequently, they evaluated their system on the BGPmon-reported potential hijacking events during August 2015 and reported a 72% detection rate with 8.2% false positives. However, scalability issues limited the adoption of their graph database filtering system.

E. Machine learning-based prefix and origin changes hijacking detection methods

The authors of [70] sought to maximize the performance of various ML algorithms and signature-based methods for hijack detection. Their signature-based approach used binary search and data reduction to narrow the search space for BGP messages, extract essential features, and feed them to ML classifiers to distinguish hijacks from normal BGP operations. They claimed 81% accuracy with J48 and Random Forest. However, their approach still suffers from 19% of false negatives and positives.

Cho et al. [13] applied the concept of AS hegemony, combined with heuristic methods, to classify BGP hijacks into pre-pending mistakes, origin changes, typos, and AS path forging. Similar to our method, the authors used a manually verified ground-truth dataset to train a Random Forest model to classify the potential hijacking events of BGPMon and BGPStream. Subsequently, 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.

The authors of [12] utilized the concept of ASN embedding to propose a deep learning-based BGP hijacking detection system. They used 19 routeviews collectors to extract ≈3.6M BGP paths, 63K AS vertices, and 113K undirected links. Moreover, the authors trained an end-to-end deep learning model based on labeled BGP routes: about 2.6M (regular) and 48K (malicious), and ≈900K routes with missing ToRs. They validated their findings based on the ground truth of Cho et al. [13] and recorded an accuracy of 99.99% with 0.00% false alarm rates. Furthermore, they tested their model on 48 hijacking events recorded between 2008 and 2018, correctly classifying 32 of them.

The authors of [71] proposed an unsupervised learning system for BGP hijacking detection that embeds both ASNs and IP prefixes using AS-paths derived from BGP updates. The authors extracted approximately 3.6M BGP paths from 19 collectors, and trained a deep learning model on a large, labeled dataset. Their dataset consists of about 2.6M regular routes and 48K malicious routes, with additional routes used for unsupervised embedding. Their system detects hijacking by identifying short-term deviations in the AS-path structure at the route level. The authors validate their results on large-scale labeled data and 48 historical hijacking events, and report a high detection rate with few false positives.

Hoarau et al. [72] present a graph neural network–based technique that uses short sequences of BGP graphs derived from BGP updates to detect large-scale event-driven BGP anomalies such as outages and route leaks. Their model leverages graph convolutional layers to capture structural changes in the AS-level topology. The authors reported an accuracy of 90% after 6 minutes and a maximum of 96%, and suggested further exploration of their adopted approach.

Wang et al. [73] propose a large-scale BGP routing anomaly detection framework based on graph attention autoencoders and a multi-layer perceptron classifier to detect prefix hijacks and route leaks. They used BGP updates to construct time-sequenced AS-level topology graphs and learn AS embeddings via unsupervised pre-training on normal routing behavior. An attention mechanism assigns higher weights to structurally important ASes, mitigating interference from less relevant neighbors. The authors evaluated their approach on several historical large-scale events, demonstrating improved accuracy compared to prior graph-based and statistical approaches.

While [71] and [73] target event-level detection based on short-term path and topology deviations, respectively, our work focuses on longitudinal AS-level behavioral patterns to

characterize serial hijackers over extended time spans.

Thomas et al. [14] recently presented DFOH, a technique that detects forged-origin hijacks, in which an attacker manipulates a potentially legitimate AS path to evade current detection methods. The authors first identified the critical inference challenges and systematically derived features to address topological biases during model training. Subsequently, they achieved 90.9% forged-origin detection in about 5 minutes, with manageable numbers of reported cases for manual verification. Although this study effectively detects forged-origin hijacks, it would miss prefix-origin hijacks with no new links, which was the primary focus of the original study [8], and ours [28].

Contrary to most generic prefix hijacking detection and mitigation approaches discussed above, Testart et al. [8] in 2019 leveraged BGP data to investigate the long-term behavioral patterns of a distinct group of ASes named serial hijackers. The authors leveraged multiple datasets to examine the dominant characteristics of serial hijackers and derived 52 statistical features to train an ExtraTrees classifier using various sampling methods. They used various supplementary data to cross-validate the flagged networks.

In conclusion, our study builds upon the foundational insights of Testart et al. [8] to advance our understanding of long-term announcement patterns of ASes to identify potential serial prefix origin hijackers. While Testart et al. relied on extensive 5-year BGP datasets to profile hijacking behaviors, our approach achieves comparable results using only 0.35% of the data. Similar to [49], we have utilized RPKI validation data (ROAs) and other supplementary datasets to quantify the proportion of potentially hijacked prefixes, thereby providing new insights into our flagged networks. Furthermore, we shed light on the ASes propagating the invalid announcements originating from our flagged ASes, including MANRS participants. Lastly, building on [7], [63], [68], we took steps to further distinguish between intentional and unintentional hijacks.

Table IV summarizes key methodological dimensions, including detection targets, input data, feature design, ground-truth sources, and sizes across prior BGP hijacking detection approaches to contextualize our approach within the broader landscape of ML-based BGP hijacking detection systems. However, it is crucial to note that this table is not a direct comparison of our approach with the listed approaches because they target different classes of BGP anomalies, such as forged-origin BGP hijacks [14], event-level detection [13], and route-level hijacking detection [12], and BGP anomaly classifications (direct, indirect, and link failure) [75]. In contrast, both the original study [8] and ours [28] focus on long-term behavioural detection of serial hijackers at the AS level, which is fundamentally different from event-oriented hijack detection. Therefore, the table aims to clarify that our contribution addresses a dedicated, previously underexplored scope of BGP hijacking detection.

In addition to the BGP data used across all approaches, both the original study [8] and ours [28] incorporate auxiliary RIR, MOAS, and bogon datasets to capture richer AS-level attributes, thereby improving the reliability of the classification task. Unlike the original study [8] and the remaining approaches [12]–[14], we further integrate RPKI validation results, MANRS membership information, and longitudinal AS-path observations to refine the interpretation of why certain ASes are flagged as potentially malicious, enabling us to distinguish potential routing inconsistencies from misconfigurations or operational anomalies. Collectively, this expanded data foundation strengthens the semantic grounding of our classifier's decisions and provides a more informed context for analyzing the behavior of ASes identified as serial hijackers.

Similar to the 70 manually curated event-level hijacking ground-truth events used in [12], [13] for training their classifier, our work relies on a fundamentally similar form of ground truth. We use 23 manually curated serial hijackers extracted from long-running NANOG incident reports. For instance, ASes that were repeatedly implicated in sustained, multi-day prefix hijacking activity. We additionally incorporate 217 MANRS-compliant operators to represent legitimate behaviour, whose routing practices undergo continuous policy and operational scrutiny. This combination yields a ground-truth dataset that reflects persistent AS-level behaviour rather than isolated routing events. In contrast, the remaining studies rely on individual anomalies, such as origin changes, path forgeries, or link failures, which capture short-lived control-plane deviations but do not characterize long-term malicious conduct. This distinction is crucial because serial-hijacker detection requires labels tied to behavioural patterns over an extended period, making event-based datasets unsuitable for training or validating classifiers targeting long-term malicious behaviour.

The model choice and feature design in both our study and Testart et al. differ from those in the remaining studies due to the structure and variety of our input data. For instance, we used ExtraTrees because it handles mixed feature types and is independent of strict feature scaling, enabling it to capture the essential attributes derived from BGP visibility, MOAS announcements, and RIR allocations. These differences arise from the specific datasets each study adopts and the representational choices required by those datasets, rather than from comparable design decisions. It is crucial to note that, although our ground truth and feature design emphasize long-term behavioural patterns, we also include several features such as the variations in daily visibility and abrupt changes in announcement volatility that capture short-term irregularities due to individual routing incidents to enable our model to remain sensitive to event-driven anomalies, a shared attribute with the detection systems in [12], [13], [75] in Table III.

There are some relevant conceptual similarities across all the approaches listed in Table III. Despite differences in temporal resolution, several methods rely on MOAS-related features and prefix-activity patterns to identify hijacking events. These shared attributes indicate that methods could benefit from one another. For instance, short-lived dynamics could refine the long-term behavioural baseline. At the same time, the longitudinal stability measures could help event-based systems identify persistent malicious behaviour from isolated misconfigurations.

Finally, unlike the feature design approaches in [12]–[14], which incorporate AS-path structural features to detect AS path manipulation behaviour, both Testaart et al. and our study

TABLE IV
COMPARISON OF ML-BASED BGP HIJACKING DETECTION SYSTEMS' DETECTION SCOPE, DATA REQUIREMENTS, GROUND TRUTH, AND METHODOLOGY

Study	Detection Target	Input Data	Ground Truth Source	Ground T. Size	ML Algorithm	Main Features
Our study [28]	Serial BGP hijackers	Five-minute daily snapshots of peer-prefix-origin bogons; MOAS, RIR	NANOG reports and MANRS members	217 normal and 23 malicious ASes	Extra Trees	44 behavioral features: visibility; prefix stability; longevity; fragmentation; MOAS
Testart et al. [8]	Serial BGP hijackers	288 daily snapshots of peer-prefix-origin bogons; MOAS, RIR	NANOG reports + MANRS members	217 normal and 23 malicious ASes	Extra Trees	52 behavioral features: visibility; prefix stability; longevity; fragmentation; MOAS
Cho et al. [13]	Origin changes, typos, prepending	BGPmon alarms and Dyn blog incidents	Dyn blog and BGPmon reports	70 events	Random Forest	AS hegemony; Edit distance; Prepending mistakes; Binary MOAS
Shapira [12]	Route-level BGP hijack detection	RouteViews's BGP paths	VF-labelling and manual auditing	2.7M GREEN 48.8K RED routes	LSTM	ASN embeddings (BGP2Vec)
Shapira & Shavitt [71]	IP prefix hijack detection	RouteViews's BGP paths	Dyn blog, MANRS and BGPmon reports	13 BGP hostorical events	Doc2Vec (PV-DM [74])	IP Prefix embeddings (AP2Vec)
Paiva et al. [75]	BGP anomaly types: direct, indirect, link failures	BGP updates (1-min batches)	Historical events (e.g., Nimda, AWS Leak Slammer)	9 BGP hostorical events	LSTM	AS relationship; AS geo. location; anamoulous event attributes
Holterbach et al. [14]	Forged-origin BGP hijacks	AS paths of BGP updates from RIPE RIS and RouteViews	Simulations and historical BGP hijacking events	Not Specified	Random Forest	Topological; peering; AS-path patterns and bidirectionality features

rely solely on announcement behavioural patterns, making our approaches unsuitable for detecting path manipulation. Therefore, these AS-path-based signals provide complementary information that could broaden our detection scope to include path-manipulation hijacks. At the same time, our long-term stability features could help these systems distinguish sustained malicious behaviour from isolated or short-lived anomalies.

XI. DISCUSSION

A. Responding to hijacking

The operators, researchers, and policymakers need to decide how to address serial hijackers. For instance, Section IX shows that alleged serial hijackers often remain active or at least allocated by the RIR.

One way forward could be a reputation system that classifies ASes based on their prefix-origination behavior and makes their reputation scores public. Additionally, not peering with ASes of low reputation could incentivize 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. Similar to [76], the "naming and shaming" approach could improve the security posture of the interdomain routing ecosystem, thereby promoting transparency and accountability.

A reputation system could enhance the security posture and resilience of the Interdomain routing ecosystem. However, there could be legal issues arising from incorrectly flagging benign networks (false positives), which could damage a network's reputation and negatively affect its chances of competing in the open market. Therefore, we recommend a collaborative effort among academic researchers, industry, legal practitioners, and all stakeholders involved in Internet governance to develop standard frameworks, evaluation metrics, and testbeds to enable rigorous evaluation and approval of research products 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 reputations of wrongly flagged networks, ensure fair competition in an open market, and enhance the Internet's resilience and security.

B. Immediate actionable recommendation

Our findings show that there is still some room for improvement despite the increased adoption of RPKI's Route Origin Authorization (ROA) and Route Origin Validation (ROV).

Therefore, we recommend that network operators register ROAs for all their prefixes, including the unused ones, and enforce route validation to enable a more resilient and secure interdomain routing ecosystem. This could be further stimulated using collaborative efforts and extensive monitoring.

a) The need for a cross-operator alliance: Our upstream analysis of the RPKI-invalid announcements shows that a small fraction of the global routing ASes propagate these prefixes from our flagged ASes, as discussed in Section VIII-D, including some MANRS participants (see Section VIII-E). The systemic vulnerabilities of a single AS that propagates an invalid announcement can undermine the integrity and security of inter-domain routing. Therefore, we recommend a cross-operator alliance among large network operators, especially the few identified potentially vulnerable ASes and large carriers,

to enforce RPKI-ROV on all inbound routes to ensure only valid routes are accepted and propagated.

b) A holistic approach to enforce MANRS Actions: Additionally, the 2022 KLAYSWAP hijacking events revealed that many large transit networks, including MANRS participants such as Zayo, inadvertently propagate hijacked routes [18], thereby violating MANRS conformance and compromising the security and resilience of the Internet. This security breach shows that being a MANRS participant alone cannot ensure resilient, secure traffic routing, especially when robust validation of incoming routes is not implemented and enforced. As a result, we advocate regular, effective audits of MANRS participants' compliance with routing security BCPs, such as filtering.

C. Trade-off analysis and limitations

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

a) Short lived hijacks: Contrary to the original study [8], our single-daily snapshot is prone to missing short-lived hijacks, such as potentially malicious events that start and end within the same day. The trade-off of choosing fewer snapshots and missing short-lived hijacks remains an open challenge. For example, consistent with the findings of [41], [45], our exploratory analysis shows that our peer-pfx-origins dataset contains significant redundancy, which could distort our analysis. Therefore, an adaptive sampling method could help us select a subset of the data to capture crucial events and streamline our analysis to develop efficient detection methods. AS a result, enabling us to conduct such resource-intensive research with minimal resources. This adaptive approach could increase the sampling rate during periods with many BGP control-plane events and lower it during periods with fewer events.

The single daily snapshots have no impact on our findings and conclusions, as our study's primary focus is the long-term behavioural patterns of serial hijackers, which mostly lasted more 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 indicate that our sampling methods should capture the dominant attributes of serial hijackers, thereby minimizing bias in our analysis and understanding of their long-term behavioural patterns. Although our main objective and conclusions apply to the long-term behavioural patterns of serial hijackers, detecting short-lived serial hijackers and understanding how such events could affect our sampling strategies remain open research challenges.

Similar to the original study [8], our study is limited to the prefix-origin-based behavioural pattern profiling and cannot detect AS path manipulations hijacking events. Holterbach et al. [14] have attempted to address this security vulnerability with a system called Detect Forged-Origin Hijack (DFOH) that uses various machine-learning features to infer potential suspicious AS links. Therefore, developing a machine learning classifier that combines our approach with AS-path-level inference is an interesting research direction for future work.

b) False positives: Moreover, our classifier recorded false positives, such as flagging DDoS protection services as serial hijackers, misconfigurations by upstream providers and their customers, and lease IP address space (see Section VIII-F).

As a result, we advocate that network operators of provider and customer networks, and IP leasing companies, register or update their ROA records promptly to reflect the new operational reality. This would significantly reduce the false positives of existing systems and ensure accurate detection of potential malicious events.

In the future, we also intend to use Recurrent Neural Networks (RNNs) and word embeddings, such as Word2Vec, to extract crucial features that capture semantic meaning. For instance, the authors of [12] extended the Word2Vec concept to BGP2Vec, which uses AS-paths to embed each ASN into a dense vector and applies it to an RNN to classify BGP hijacks. The authors claimed an accuracy of 99.99% and 0.00% false positives, which could help minimize false positives in our classifier.

c) Reproducibility challenges: Lastly, although we found the original paper's methodology and findings to be as claimed, reproducing them was nontrivial due to numerous missing artifacts that we had to figure out ourselves. Therefore, publishing papers with supporting artifacts should be encouraged. As a result, we are making all our artifacts, including our data collection and preprocessing scripts, available to improve transparency and complement research efforts.[6]

D. Ethical considerations

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

XII. 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 using much less data and have developed a competing classifier that evaluates current trends in serial hijacking activity. 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 finding 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, indicating that the routing community has not found a common path to address malicious behavior. Additionally, we have shown how widely RPKI-invalid announced prefixes propagate on the Internet, which remains a systemic vulnerability to the routing

[6]https://github.com/ebrimajaw/revisiting-serial-BGP-hijackers

infrastructure. Moreover, we have determined the few ASes on the AS paths propagating these RPKI-invalids, which include a few MANRS participants. These results highlight the need for a holistic approach to routing security. 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 for their constructive and valuable feedback. 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] C. K. Sanders and E. Scanlon, "The digital divide is a human rights issue: Advancing social inclusion through social work advocacy," *Journal of Human Rights and Social Work*, vol. 6, no. 2, pp. 130–143, 2021. [Online]. Available: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7973804/

[2] M. Cheng, Q. Xu, J. L.V., W. Liu, Q. Li, and J. Wang, "MS-LSTM: A multi-scale LSTM model for BGP anomaly detection," in *2016 IEEE 24th International Conference on Network Protocols (ICNP)*. Singapore: IEEE, Nov. 2016, pp. 1–6. [Online]. Available: https://ieeexplore.ieee.org/document/7785326/

[3] Q. Tang, F. R. Yu, R. Xie, A. Boukerche, T. Huang, and Y. Liu, "Internet of Intelligence: A Survey on the Enabling Technologies, Applications, and Challenges," *IEEE Communications Surveys & Tutorials*, vol. 24, no. 3, pp. 1394–1434, 2022. [Online]. Available: https://ieeexplore.ieee.org/document/9775170/

[4] 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/

[5] 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.

[6] L. Miller and C. Pelsser, "A Taxonomy of Attacks Using BGP Blackholing," in *Computer Security – ESORICS 2019*, K. Sako, S. Schneider, and P. Y. A. Ryan, Eds., vol. 11735. Cham: Springer International Publishing, 2019, pp. 107–127, series Title: Lecture Notes in Computer Science. [Online]. Available: http://link.springer.com/10.1007/978-3-030-29959-0_6

[7] J. Schlamp, G. Carle, and E. W. Biersack, "A Forensic Case Study on AS Hijacking: The Attacker's Perspective," *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 2, p. 6, 2013.

[8] 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

[9] 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/

[10] R. S. Silva and L. M. Felipe de Moraes, "GonoGo - assessing the confidence level of distribute intrusion detection systems alarms based on BGP," *IEEE Transactions on Network and Service Management*, vol. 22, no. 1, pp. 209–219, 2025. [Online]. Available: https://ieeexplore.ieee.org/document/10695796/

[11] A. H. Feng, P. Francois, M. Younsi, S. Frénot, T. Graf, W. Du, P. Lucente, and A. Elhassany, "Detecting service disruptions in large BGP/MPLS VPN networks," *IEEE Transactions on Network and Service Management*, vol. 22, no. 5, pp. 3964–3977, 2025. [Online]. Available: https://ieeexplore.ieee.org/document/11078445

[12] 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

[13] 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/

[14] T. Holterbach, T. Alfroy, A. Phokeer, A. Dainotti, and C. Pelsser, "A system to detect forged-origin bgp hijacks," in *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24)*, 2024, pp. 1751–1770. [Online]. Available: https://www.usenix.org/conference/nsdi24/presentation/holterbach

[15] 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 *Proceedings of the 1st ACM SIGCOMM Workshop on Internet measurement*, ser. IMW '01. New York, NY, USA: Association for Computing Machinery, Nov. 2001, pp. 31–35. [Online]. Available: https://doi.org/10.1145/505202.505207

[16] DDoS mitigation firm has history of hijacks – krebs on security. [Online]. Available: https://krebsonsecurity.com/2016/09/ddos-mitigation-firm-has-history-of-hijacks/

[17] 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/

[18] 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/

[19] 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

[20] B. Herdes, M. Zhang, and T. Ryan. Cloudflare 1.1.1.1 incident on june 27, 2024. [Online]. Available: https://blog.cloudflare.com/cloudflare-1111-incident-on-june-27-2024

[21] R. Bush and R. Austein, "The resource public key infrastructure (RPKI) to router protocol, version 1," Internet Engineering Task Force, Request for Comments RFC 8210, 2017-09, num Pages: 35. [Online]. Available: https://datatracker.ietf.org/doc/rfc8210

[22] M. Lepinski, D. Kong, and S. Kent, "A profile for route origin authorizations (ROAs)," Internet Engineering Task Force, Request for Comments RFC 6482, 2012-02, num Pages: 9. [Online]. Available: https://datatracker.ietf.org/doc/rfc6482

[23] D. Mirdita, H. Schulmann, and M. Waidner. SoK: An introspective analysis of RPKI security. [Online]. Available: https://arxiv.org/html/2408.12359v1

[24] R. Bush, "Origin validation operation based on the resource public key infrastructure (RPKI)," Internet Engineering Task Force, Request for Comments RFC 7115, 2014-01, num Pages: 11. [Online]. Available: https://datatracker.ietf.org/doc/rfc7115

[25] 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

[26] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang, "PHAS: a prefix hijack alert system," in *Proceedings of the 15th conference on USENIX Security Symposium - Volume 15*, ser. USENIX-SS'06, no. 14. USA: USENIX Association, 2006-07-31, pp. 153–166.

[27] 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/

[28] E. Jaw, M. Müller, C. Hesselman, and L. Nieuwenhuis, "Serial BGP hijackers: A reproducibility study and assessment of current dynamics," in *2024 8th Network Traffic Measurement and Analysis Conference (TMA)*, 2024-05, pp. 1–10.

[29] T. Alfroy, T. Holterbach, T. Krenc, K. C. Claffy, and C. Pelsser, "The next generation of BGP data collection platforms," in *Proceedings of the ACM SIGCOMM 2024 Conference*, ser. ACM SIGCOMM '24. New York, NY, USA: Association for Computing Machinery, 2024-08-04, pp. 794–812. [Online]. Available: https://dl.acm.org/doi/10.1145/3651890.3672251

[30] T. Krenc, R. Beverly, and G. Smaragdakis, "Keep your communities clean: exploring the routing message impact of BGP communities," in *Proceedings of the 16th International Conference on emerging Networking EXperiments and Technologies*, ser. CoNEXT '20.

Association for Computing Machinery, 2020-11-24, pp. 443–450. [Online]. Available: https://doi.org/10.1145/3386367.3432731

[31] 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/

[32] 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

[33] 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

[34] T. Hlavacek, H. Shulman, N. Vogel, and M. Waidner, "Keep your friends close, but your routeservers closer: Insights into RPKI validation in the internet," 2023-03-21. [Online]. Available: http://arxiv.org/abs/2303.11772

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

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

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

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

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

[40] 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/

[41] 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

[42] O. Darwich, C. Pelsser, and K. Vermeulen, "Detecting traffic engineering from public BGP data," in *Passive and Active Measurement*, C. Testart, R. Van Rijswijk-Deij, and B. Stiller, Eds. Springer Nature Switzerland, 2025, vol. 15567, pp. 307–334, series Title: Lecture Notes in Computer Science. [Online]. Available: https://link.springer.com/10.1007/978-3-031-85960-1_13

[43] K. Z. Sediqi, L. Prehn, and O. Gasser, "Hyper-specific prefixes: gotta enjoy the little things in interdomain routing," *SIGCOMM Comput. Commun. Rev.*, vol. 52, no. 2, pp. 20–34, 2022-06-20. [Online]. Available: https://doi.org/10.1145/3544912.3544916

[44] G. Huston. (2017) BGP more specifics: Routing vandalism or useful? [Online]. Available: https://labs.ripe.net/author/gih/bgp-more-specifics-routing-vandalism-or-useful/

[45] E. Jaw, T. Krenc, M. Müller, K. Claffy, L. Nieuwenhuis, and C. Hesselman, "Noisy neighbours: Keep the neighbourhood quiet," in *21st International Conference on Network and Service Management (CNSM)*, 2025.

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

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

[48] H. Schulmann and S. Zhao, "Learning to identify conflicts in RPKI," in *Proceedings of the 20th ACM Asia Conference on Computer and Communications Security*, ser. ASIA CCS '25. New York, NY, USA: Association for Computing Machinery, 2025, pp. 1490–1505. [Online]. Available: https://dl.acm.org/doi/10.1145/3708821.3710833

[49] M. Wählisch, O. Maennel, and T. C. Schmidt, "Towards detecting BGP route hijacking using the RPKI," *ACM SIGCOMM Computer Communication Review*, vol. 42, no. 4, pp. 103–104, Aug. 2012. [Online]. Available: https://doi.org/10.1145/2377677.2377702

[50] K. Michalak and H. Kwasnicka, "Correlation based feature selection method," *Int. J. Bio-Inspired Comput.*, vol. 2, no. 5, pp. 319–332, 2010. [Online]. Available: https://doi.org/10.1504/IJBIC.2010.036158

[51] S. Matharaarachchi, M. Domaratzki, and S. Muthukumarana, "Minimizing features while maintaining performance in data classification problems," *PeerJ Computer Science*, vol. 8, p. e1081, 2022. [Online]. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC9575878/

[52] P. Geurts, D. Ernst, and L. Wehenkel, "Extremely randomized trees," *Machine Learning*, vol. 63, no. 1, pp. 3–42, 2006. [Online]. Available: https://doi.org/10.1007/s10994-006-6226-1

[53] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001. [Online]. Available: https://doi.org/10.1023/A:1010933404324

[54] CleanTalk. (2024) ASN list - ASN lookup blacklist and IP's spam statistics. [Online]. Available: https://cleantalk.org/blacklists/asn?page=11

[55] 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

[56] ASRank. (2024) AS rank : GraphQL API overview page. [Online]. Available: https://api.asrank.caida.org/v2/docs

[57] RIPE NCC. Rpki snapshots/. [Online]. Available: https://ftp.ripe.net/rpki/

[58] E. Jaw, M. Müller, C. Hesselman, and L. Nieuwenhuis, "Poster: Quantifying the proportion of hijacked prefixes among the identified prefix hijackers," 2023-09-28.

[59] P. Marcos, L. Prehn, L. Leal, A. Dainotti, A. Feldmann, and M. Barcellos, "AS-path prepending: there is no rose without a thorn," in *Proceedings of the ACM Internet Measurement Conference*, ser. IMC '20. Association for Computing Machinery, 2020-10-27, pp. 506–520. [Online]. Available: https://doi.org/10.1145/3419394.3423642

[60] A. Milolidakis, T. Buhler, M. Chiesa, L. Vanbever, and S. Vissicchio, "Poster: Smart BGP Hijacks that Evade Public Route Collectors," in *Special Interest Group on Data Communication 2021*, undefined 2021, p. 2.

[61] B. Du, C. Testart, R. Fontugne, G. Akiwate, A. C. Snoeren, and k. claffy, "Mind your MANRS: measuring the MANRS ecosystem," in *Proceedings of the 22nd ACM Internet Measurement Conference*, ser. IMC '22. Association for Computing Machinery, 2022-10-25, pp. 716–729. [Online]. Available: https://dl.acm.org/doi/10.1145/3517745.3561419

[62] L. Qin, L. Chen, D. Li, H. Ye, and Y. Wang, "Understanding route origin validation deployment in the real world and why MANRS action 1 is not followed," in *Proceedings 2024 Network and Distributed System Security Symposium*. Internet Society, 2024. [Online]. Available: https://www.ndss-symposium.org/wp-content/uploads/2024-214-paper.pdf

[63] B. Du, R. Fontugne, C. Testart, A. C. Snoeren, and k. claffy, "Sublet your subnet: Inferring IP leasing in the wild," in *Proceedings of the 2024 ACM on Internet Measurement Conference*, ser. IMC '24. Association for Computing Machinery, 2024-11-04, pp. 328–336. [Online]. Available: https://dl.acm.org/doi/10.1145/3646547.3689010

[64] The world's first IP marketplace. trusted, secure & reliable. [Online]. Available: https://www.ipxo.com

[65] RIPE NCC, "Routing Information Service (RIS)," https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris, 2023.

[66] Z. Zhang, Y. Zhang, Y. C. Hu, Z. M. Mao, and R. Bush, "Ispy: detecting ip prefix hijacking on my own," in *Proceedings of the ACM SIGCOMM 2008 conference on Data communication*. ACM, 2008-08-17, pp. 327–338. [Online]. Available: https://dl.acm.org/doi/10.1145/1402958.1402996

[67] 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

[68] P.-A. Vervier, Q. Jacquemart, J. Schlamp, O. Thonnard, G. Carle, G. Urvoy-Keller, E. Biersack, and M. Dacier, "Malicious BGP hijacks: Appearances can be deceiving," in *2014 IEEE International Conference on Communications (ICC)*, 2014-06, pp. 884–889, ISSN: 1938-1883. [Online]. Available: https://ieeexplore.ieee.org/document/6883431

[69] 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 *Traffic Monitoring and Analysis*, M. Steiner, P. Barlet-Ros, and O. Bonaventure, Eds. Springer International Publishing, 2015, vol. 9053, pp. 173–187, series Title: Lecture Notes in Computer Science. [Online]. Available: http://link.springer.com/10.1007/978-3-319-17172-2_12

[70] H. Alshamrani and B. Ghita, "IP prefix hijack detection using BGP connectivity monitoring," in *2016 IEEE 17th International Conference on High Performance Switching and Routing (HPSR)*,

2016-06, pp. 35–41, ISSN: 2325-5609. [Online]. Available: https://ieeexplore.ieee.org/document/7525636
[71] T. Shapira and Y. Shavitt, "AP2vec: An unsupervised approach for BGP hijacking detection," *IEEE Transactions on Network and Service Management*, vol. 19, no. 3, pp. 2255–2268, 2022. [Online]. Available: https://ieeexplore.ieee.org/document/9754706
[72] K. Hoarau, P. U. Tournoux, and T. Razafindralambo, "BGNN: Detection of BGP anomalies using graph neural networks," in *2022 IEEE Symposium on Computers and Communications (ISCC)*, 2022, pp. 1–6, ISSN: 2642-7389. [Online]. Available: https://ieeexplore.ieee.org/document/9912989
[73] Z. Wang, F. Yuan, H. Qiu, Y. Liu, and X. Luo, "Large-scale BGP routing anomaly detection based on graph attention auto-encoder," *IEEE Transactions on Network and Service Management*, vol. 23, pp. 487–501, 2026. [Online]. Available: https://ieeexplore.ieee.org/document/11328038
[74] Q. Le and T. Mikolov, "Distributed representations of sentences and documents," in *Proceedings of the 31st International Conference on Machine Learning*. PMLR, 2014, pp. 1188–1196, ISSN: 1938-7228. [Online]. Available: https://proceedings.mlr.press/v32/le14.html
[75] T. B. Paiva, Y. Siqueira, D. M. Batista, R. Hirata, and R. Terada, "BGP anomalies classification using features based on AS relationship graphs," in *2021 IEEE Latin-American Conference on Communications (LATINCOM)*, 2021, pp. 1–6, ISSN: 2330-989X. [Online]. Available: https://ieeexplore.ieee.org/document/9647824
[76] 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.



Ebrima Jaw has been a PhD student in the Design and Analysis of Communication Systems (DACS) at the University of Twente, The Netherlands, since 2022. He received his BSc in Information Systems from the University of Gambia and his MSc in Computer Science and Technology (Network and Information Security) in 2022 from Guizhou University, China. His current research interests include interdomain routing security, Internet measurement, and the scalability of Internet infrastructure.



Moritz Müller is a research engineer at SIDN Labs, the research department of the operator of the Dutch ccTLD. He is also a guest researcher at the University of Twente. In his research, Moritz is working on the security and reliability of core components of the Internet, focusing mainly on the Border Gateway Protocol (BGP) and the Domain Name System (DNS). He is also a co-chair of the DNS working group at RIPE.



Cristian Hesselman leads the Trusted Digital Infrastructures unit at SURF, overseeing the core network, security, and identity services for the Dutch research and education community. He is a professor at the University of Twente and a member of the Netherlands' Cybersecurity Council. Before joining SURF in early 2026, Cristian spent over 14 years directing SIDN Labs, the research arm for the .nl top-level domain. His research focuses on enhancing the trustworthiness of the Internet and future networks through large-scale infrastructure measurements and the development of new Internet security mechanisms.



Lambert J.M. Nieuwenhuis is a full professor at the Department of HighTech Business and Entrepreneurship of the Behavioural, Management and Social Sciences Faculty at the University of Twente. He is chair of the Supervisory Board of NLnet Labs, developing open source software for DNS and BGP routing. He received a BSc (1977) and MSc (1980) in Electrical Engineering and a PhD (1991) in Computer Science. He combined a more than thirty-year academic career at various universities with international R&D in the telecommunications industry and independent consultancy in innovation management. His current research interests include Service Innovation, Cyber Security, and Internet Economics.