Metis: Detecting Fake AS-PATHs based on Link Prediction

Chengwan Zhang*, Congcong Miao*, Changqing An*, Anlun Hong*,
Ning Wang†, Zhiquan Wang* and Jilong Wang*
*Institute for Network Sciences and Cyberspace, Tsinghua University, China
†Institute for Communication Systems, University of Surrey, UK

***Abstract*—BGP route hijacking is a critical threat to the Internet. Existing works on path hijacking detection firstly monitor the routes of the whole network and then directly trigger a suspicious alarm if the link has not been seen before. However, these naive approaches will cause false positive identification and introduce unnecessary verification overhead. In this work, we propose Metis, a matching-and-prediction system to filter out normal unseen links. We first use a matching method with three rules to find out suspicious links if there is an unseen AS. Otherwise, we propose using a neural network to make a prediction based on the AS information at each end of the link and further quantify the suspicion level. Our large-scale simulation results show that Metis can achieve precision and recall of over 80% for detecting fake AS-PATHs. Moreover, our deployment experiences show that compared to state-of-the-art system, Metis can save 80% overhead.**

***Index Terms*—BGP anomaly, Prefix hijack detection, Link prediction**

I. Introduction

Autonomous systems (ASes) use border gateway protocol (BGP) [1] to advertise a set of IP prefixes and establish inter-domain routes in the Internet. The current Internet has almost 80,000 ASes to exchange route reachability information with each other. However, the BGP protocol lacks route authenticity verification mechanism, which allows attackers to inject bogus routes into the routing system for the purpose of redirecting traffic. This is known as BGP route hijacking or prefix hijacking. Prefix hijacking can take the form of *origin hijacking* or *path hijacking*. In origin hijacking, the attacker simply acts as the origin AS to announce IP prefix belongs to another AS. However, this triggers a conflict known as Multiple Origin Autonomous System (MOAS) [2], which can be detected by existing systems such as [3]–[5]. To circumvent MOAS, some attackers announce IP prefixes while tampering with the AS-PATH, which is known as path hijacking. Early prefix hijackings were mainly origin hijackings caused by misconfiguration, such as the Pakistan Telecom hijacking YouTube in 2008 [6] and the China Telecom hijacking event in 2010 [7]. However, over the past few years, path hijacking has emerged as a favored method for attackers as it is a stealthier form of attack, as evidenced by incidents such as [8], [9].

Origin hijacking has been well studied [2], [10], but there are fewer works on path hijacking [3], [11], [12], which is more stealthy and harder to detect. In the control plane, one important indication of path hijacking is the appearance of unseen links, which are typically introduced by attackers who manipulate the AS-PATH (§II.A). However, existing state-of-the-art methods (such as Fingerprints [12], Argus [3], and Artemis [11]) trigger a suspicious alarm whenever a new link that has not been seen in the past period is detected. Since many new links are emerging every day, but only a few of them are abnormal, these methods tend to generate numerous false alarms or encounter unnecessary data plane overhead for verifying normal unseen links.

To address this problem, we propose a method that uses link prediction to evaluate the authenticity of unseen links, thereby filtering out normal unseen links. Specifically, the Internet can be modeled as an AS-level network (AS topology) that continuously evolves, with links added or removed every day. Intuitively, links of AS topology are not generated randomly because many factors (such as country, geographical location, type, size, and tier of the ASes) can influence the interconnection of ASes. For example, a small stub AS may prefer to connect to a local ISP to save money on network cables, whereas a global transit AS may try to connect with ASes everywhere and gain as many customers and peers as possible. The non-randomness of link generation in AS topology leads to its regularity, enabling the evaluation of the authenticity of unseen links based on seen links. This is where link prediction comes into play. Link prediction is a task that calculates the probability of an unknown link's existence based on the observed links. In this paper, we propose Metis, a link prediction based real-time fake AS-PATHs detection system. Metis first use a matching method with three rules to find out suspicious links if there is an unseen AS. Otherwise, we classify the rest unseen links with a GNN-based link predictor and further quantify the suspicion level with the characteristics of the fake AS-PATH.

The followings are the main contributions of this paper:

- 1. We propose to evaluate authenticity of unseen AS links with link prediction. We use a Graph Neural Network (GNN)-based link predictor to classify unseen links, achieving up to an accuracy of 95% and AUC of 0.98.
- 2. We propose a real-time fake AS-PATH detection system Metis by combining link prediction and rules, which achieve precision and recall of over 80% for detecting fake AS-PATHs, and can save the-state-of-the-art more than 80% cost.

979-8-3503-0048-2/23/$31.00 ©2023 IEEE

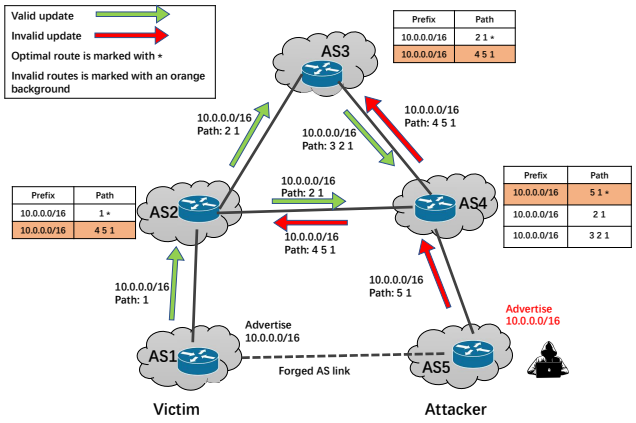

Fig. 1. Path hijack and unseen fake link

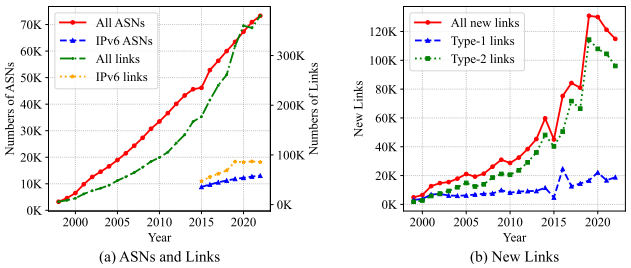

Fig. 2. Evolving of AS topology

II. Background and Motivation

A. Path hijacking and unseen link

BGP has no route authenticity verification mechanism, which is exploited by attackers to launch prefix hijacking attacks. Prefix hijacking can be categorized as origin hijacking and path hijacking. In origin hijacking, attacker AS simply advertises IP prefix of victim AS. Origin hijacking triggers MOAS conflict which can be easily detected by MOAS-based hijacking detection systems [4], [5], [13], [14]. To bypass MOAS, attackers manipulate the AS-PATH, a technique known as path hijacking. Unfortunately, path hijacking usually introduces unseen fake AS links in the AS-PATH. In Figure I, AS5 launches a path hijacking by adding AS1 to the end of AS-PATH. Leading AS2-4 to see an unseen link "AS5-AS1". Current state-of-the-art techniques [3], [11], [12] detect path hijacking based on unseen new links. These systems either directly treat the presence of new links in the past period as anomalies or mark them as suspicious and then verify them with the data plane probing. However, as the Internet sees many new links daily, many of which are not related to path hijacking, false positives are a concern.

Figure 2(a) [1] shows the evolution of the AS topology. The number of observable ASes increases linearly from 1998 to 2022 at a speed of approximately 3K ASes per year, reaching 73,303 ASes by 2022. Meanwhile, the number of observable AS links grows as a curve and reaches 379,981 in 2022. Figure 2(b) illustrates the newly emergent AS links each year, classified into two types: *Type-1 links* and *Type-2 links*. Type-1 links refer to new links with unseen ASes, while Type-2 links are newly emerging links of existing ASes. The figure indicates that AS topology generates a growing number of new links every year, with Type-2 links being the most common. In 2022, approximately 314 new links appeared daily, and 83.6% of them were Type-2 links. These new links occur due to various reasons besides path hijacking such as new BGP peering, backup links, route policy changing, BGP poisoning [16] etc. Obviously, most of them are real. If all unseen links are treated directly as anomalies, there will be many false positives. This motivates us to find a way to filter the normal unseen links.

[1]The CAIDA AS relationship dataset [15] was used to obtain all topology data used in this paper.

B. Link prediction

Link prediction is to infer whether two nodes in a network are likely to have a link. Lu et al. [17]. define link prediction as follows: for an observed simple acyclic graph or network $G(V, E)$, where V is the set of nodes and E is the set of edges. U denotes the full set of edges of size $\frac{n\cdot(n-1)}{2}$ where n is the number of nodes V. $U - E$ is the set of unobserved edges. It assumes there are some missing links (or the links that will appear in the future) in the set $U - E$, and the task of link prediction is to find out these links. Link prediction can output a value between 0 and 1 for a link in $U - E$ based observed links E, which represents how likely the link exists (or will show up in the near future). Our basic idea is to formulate the task of evaluating the authenticity of an unseen link to a link prediction problem.

III. Metis: A Real-time Link Prediction based Fake AS-PATHs Detection Framework

In this section, we will detail Metis, a real-time link prediction based fake AS-PATHs detection framework.

A. Overview

Figure 3 shows the overview of Metis. There are four key components in Metis: **Reliable AS link library (§III.B), link predictor (§III.C), Type-1 rules (§III.D) and Type-2 rules (§III.F)**.

As Figure 3 shows, Metis works as follows. ❶First, it retrieves BGP UPDATE messages from real-time BGP feed (like CAIDA bgpstream API [18]) and extracts the AS links from AS-PATH. ❷Next, it checks whether the AS link is in the reliable AS link library. If yes, then outputs "valid" directly.❸ if not, it then checks whether the AS link contains unseen AS (not in the reliable link library). We call these links *Type-1 links*. If yes, it then uses Type-1 rules to determine if it is valid or not. ❹If the link is not a Type-1 link, then the AS link is a *Type-2 link*. it will be put into the well-trained link predictor and then output a *prediction value* between 0 to 1, which means the probability that the link is valid. If the prediction value is greater than the preset threshold, it is considered a valid link. Otherwise, it is an invalid link, and then Metis uses the Type-2 rules to get the suspicion level for the invalid Type-2 link and its AS-PATH.

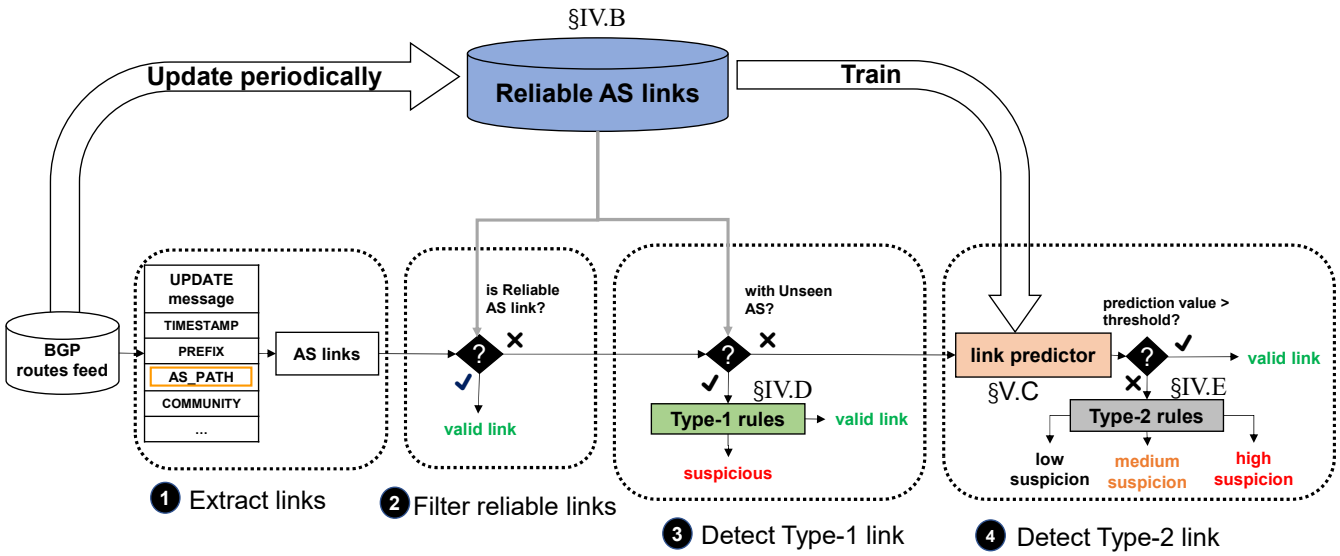

Fig. 3. Overview of Metis

If all AS links of an AS-PATH are valid, the AS-PATH is tagged valid. If not, the AS-PATH is tagged suspicious and outputs an event with suspicious links and matched rules.

B. Reliable AS link library constructing

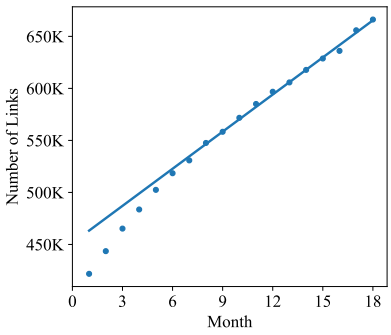

Fig. 4. The number of union AS links in CAIDA AS relationship data of the past N months of November 2021

The reliable AS link library is the foundation of Metis, in which the links are believed to be real links on the current AS topology. We calculate the union of the past N months of the CAIDA AS relationship dataset as the reliable AS link library. To determine the appropriate value of N, we count the number of union AS links in CAIDA data of the past N months of November 2021 (the time we first do the experiments). As 4 shows, the numbers of AS links grow fast at the beginning, which is due to the addition of a large number of backup links, then the growth rate gradually decreases because there are fewer and fewer undiscovered backup links. Finally, the growth rate becomes stable, which means that the links added to the union are out-of-date links because these links don't appear in any of the later RIBs again. The knee point of the curve is at N=6, which means a link not appearing in the past six months is rarely likely to come back again. So we take the value of N as 6 as our experiment setting.

Note that the above is the initial AS link library constructing due to a cold start, and there will inevitably be a few fake links in it. The reliable AS link library should be updated periodically (e.g., once a month), and the links detected as invalid can not be included.

C. Link Predictor

There are many algorithms available for link prediction, but their effectiveness can vary depending on the AS topology. Therefore, selecting the right algorithm is critical. Our chosen algorithm is SEAL [19] for two reasons: (1) SEAL is mathematically proven to be equivalent to any similarity-based heuristic algorithm; and (2) SEAL can also learn from latent and explicit features. Latent features are low-dimensional vector representations of nodes obtained from matrix factorization or network embedding methods, which focus on the global structural features of the graph. Explicit features, such as the location, country, and size of the AS in the AS topology, are side information of the nodes that are not related to the topology structure of the network.

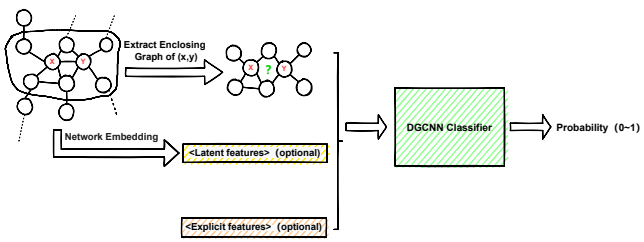

Fig. 5. Workflow of SEAL framework. For each target link, SEAL extracts a local enclosing subgraph around it, and uses the DGCNN classifier to learn general graph structure features for link prediction.

The workflow of SEAL is shown in Figure 5. SEAL [19] models the link prediction problem as a classification prob-

lem and uses a Deep Graph Convolutional Neural Network (DGCNN) [20] as its internal model. The DGCNN receives the *k-hop enclosing subgraph* of the target node pair and optional node features as input and outputs a probability of the existence of a link between the node pair. Where the h-hop closed subgraph is defined as follows: for a graph $G(V; E)$, given a node pair (x, y), the h-hop closed subgraph of (x, y) is the subgraph $G(x, y)$ formed by the set of nodes $\{i|d(i, x) \leq h \text{ or } d(i, y) \leq h\}$, where $d(x, y)$ is the distance between node x and node y. Notably, the DGCNN is trained using both positive samples (node pairs with a link) and negative samples (node pairs with no link). We did unseen link classification experiments with SEAL on the CAIDA AS topology in 2022, and the classification accuracy and AUC reached 0.95 and 0.98, respectively.[2]

D. Type-1 Rules

Type-1 links contain ASN that is not in the reliable AS link library, which are not able to be predicted by the link predictor. These newly emerging ASes may appear due to normal BGP operations, misconfigurations and BGP hijackings. We use the following rules for checking invalid Type-1 links:

- The link contains an AS that cannot be found in the AS registration records of the five RIRs. Typically, an ASN should be registered in the RIR before it can be used on the Internet. The unregistered ASNs could be caused by misconfiguration and malicious hijacking.
- The link contains an ASN that should not appear on the Internet defined in [21]. Administrators may forget to filter private-used ASNs (64512-65534, 4200000000-4294967294), resulting in them leaking to the Internet.
- The new AS in the link is not the last hop in the AS-PATH. Our experiments in §IV.D show that 97% of newly registered ASes appear on the Internet as a stub AS, and the rest are private-used ASNs, misconfigurations, IXP ASNs, and ASNs transiting for its sibling.

A Type-1 link (with its AS-PATH) that matches any of the above rules is considered suspicious, otherwise considered valid.

E. Type-2 Rules

The link with a prediction value less than the pre-set threshold is output as an suspicious link. To further distinguish the causes and the level of suspicion, We consider the characteristics of fake AS-PATHs and propose the following rules. (Experiments in §V will demonstrate the rationality of our rules)

For each AS-PATH containing the suspicious link(s), we first set an initial suspicion score of 0. if any of the following rules is successfully matched, increase the score by 1.

- The number of unique ASes in AS-PATH is greater than the pre-set threshold. In path hijacking, the attacker needs to insert ASNs into the AS-PATH, which may result in a larger number of unique ASes in the AS-PATH than in the regular AS-PATH. We suggest the threshold is set as $1.5 \sim 2$ average AS-PATH length (approximately 4.5 hops in November 2021) in the global BGP routing table.
- The suspicious link with a single-digit ASN at the end of the AS-PATH. This rule takes into account the fact that operators can accidentally introduce single-digit ASNs at the end of the AS-PATH when performing AS-PATH Prepending (ASPP) operations.
- The Damerau-Levenshtein edit distance of the two ASNs of the suspicious link is no more than 1. Same as the previous one, misconfiguration may cause such AS links.
- The AS-PATH has AS loop, and the link is in the loop. AS loops are often caused by BGP poisoning, and the links in the loop are usually fake.
- The AS-PATH violates the valley-free rule [22]. Some Malicious BGP hijackings may insert ASNs into AS-PATH to make it violate the valley-free rule. Since the AS-PATH exists unseen Type-2 links, it is not realistic to infer the relationship of the two ASe every time an unseen link appears. So we use the "global hegemony valley" proposed by [23] to determine whether the AS-PATH has a valley.
- The AS-PATH causes traffic detour, i.e., traffic goes out of a country and then comes back to that country. Traffic detour is rarely seen (6% of AS-PATHs have detour in the global BGP routing table), and path hijacking usually causes detour.

If any of the following rules are successfully matched, the score is reduced by 2.

- The suspicious link is at the end of the AS-PATH, and the link is a domestic link (the two ASes belong to the same country). This situation often occurs when a stub AS connects with a domestic ISP. Our experiment shows that Type-2 links that meet this rule are very likely to be false alarms.

Finally, the suspicious AS-PATH will be alarmed with an event containing the following information: {time, route, suspicious links, suspicion score, matched rules}. We regard an AS-PATH with a positive score as high suspicion, a negative score as low suspicion, and a medium suspicion if the score is zero. It is important to note that some normal AS-PATHs may also match certain rules, and these rules of thumb are used for improving the detection confidence and finding causes.

IV. Evaluation

In this section, we use crafted AS-PATHs to evaluate Metis and illustrate the improvement of Metis over existing unseen link-based methods by real experiments.

A. Ground truth dataset

Fake AS-PATHs and path hijackings are very rare on the real Internet. And validating fake AS-PATH and path hijacking in the real world is also difficult work, so there is

[2]We take the links in the AS topology as positive samples and randomly sample AS links not in the AS topology as negative samples.

almost no ground truth data. To evaluate Metis, we manually create a ground truth dataset.

This dataset includes **GREEN** samples (valid AS-PATHs) and **RED** samples (fake AS-PATHs). For GREEN samples, we randomly selected 7000 existing AS-PATHs in RRC00's RIB at 00:00, 2021-11-01, because there are very few fake AS-PATHs in the RIB. For RED AS-PATHs, We consider the three most common scenarios that generate fake AS-PATHs: misconfiguration, path hijacking, and BGP poisoning.

1) Misconfiguration.: To obtain this type of AS-PATHs, we first sample 1000 AS-PATHs from the RIB and then randomly add single-digit ASN or similar ASN (Damerau-Levenshtein edit distance equals 1 to the origin AS) to the tail of the AS-PATH. We create 1000 red samples for each of the two types of misconfigurations mentioned above and name them Type-1 misconfiguration and Type-2 misconfiguration, respectively.

2) Path hijacking: In general, the observed invalid AS-PATH of a path hijack can be divided into two segments: real and fake. AS2, for example, is attempting to hijack AS1, so it creates a fake AS-PATH "2 x 1" and advertises to its neighbor ASy (note that "x-1" is a real existing link, but 2-x is not). And then ASy continues to announce the invalid route to its neighbor AS3. Finally, AS3 observes an AS-PATH "y 2 x 1", in which the segment "y 2" is the real segment (real propagation path of the route), while "x 1" is the fake segment. In this paper, we follow the term *"Type-N hijacking"*, proposed in [11], to describe different types of path hijacking. The N is the length of the fake segment. With the definition, the example above is a typical Type-2 hijacking. Next hop attack [8] is Type-1 hijacking, and origin hijacking is Type-0 hijacking. We craft the Type-N hijacking AS-PATHs as follows: we first sample some AS-PATHs from the RIB as the real segment, and then sample another AS-PATH and extract the last N ASN(s) as the fake segment, finally join the real segment and the fake segment as the Type-N hijacking AS-PATH. Please note that we craft only one fake AS link for each AS-PATH. We create 1000 for Type-1, Type-2, and Type-3 hijacking, respectively.

3) BGP Poisoning.: To get such type of AS-PATHs, we first sample some AS-PATHs in the RIB and then randomly choose ASNs as poisoned ASNs which are inserted into the AS-PATHs. We create a term *Type-N BGP Poisoning* where N is the number of poisoned ASNs. We create 1000 for Type-1 and Type-2 BGP poisoning, respectively.

Totally, we get 7000 GREEN AS-PATHs and 7000 RED AS-PATHs with RIB on 2021-11-01 00:00:00 UTC of RIPE RRC00.

B. Experiment settings

We use CAIDA data of the past six months (June-October, 2021) to construct a reliable link library and train link predictor (SEAL). We set the AS-PATH length threshold of Type-2 rules as 9 (approximately twice the average AS-PATH length at that time). We do not set the prediction value threshold t beforehand. Instead, we will explore the effect of different thresholds t below.

C. Results

Figure 6(a) shows the distribution of prediction values for Type-2 links for various scenarios. As can be seen, the prediction values for our forged links are significantly lower than the links from the RIB. For misconfiguration and BGP poisoning, more than 75% have prediction values lower than 0.4, and links from the Type-2 Scenario are much lower than Type-1. But for path hijacking, as N grows, the prediction value of the links grows rapidly, which means the higher authenticity of the forged links (we will explain why below). For the Type-2 links of AS-PATHs from RIB, 50% prediction values are close to 1, and more than 75% prediction values larger than 0.8.

We temporarily consider all these Type-2 links in GREEN AS-PATHS as real and then calculate the precision and recall of Metis' Type-2 unseen link classification task under different prediction value threshold t. Since only 187 type2 links are obtained in GREEN AS-PATHs, and thousands of links are obtained in RED AS-PATHs, i.e., the positive and negative links are unbalanced. To make the positive and negative samples balanced, we select a total of 189 (27*7) links from all types of Type-2 links of RED AS-PATHs evenly. The result is shown in Figure 6(b). The curves of precision and recall intersect around t = 0.8. And when t=0.8, Metis can about obtain a precision of 80% and recall of 80% for detecting fake Type-2 unseen links.

When the threshold t is set to 0.8, and the detailed detection result is shown in Table I. Below we analyze the results for various samples.

1) GREEN AS-PATHs: We extract 11726 unique AS links in total, of which 11181 (95.4%) are reliable AS links, 358 (3.1%) are Type-1 links, and 187 (1.6%) are Type-2 links. There is a total of 34 AS-PATHs tagged suspicious. We checked these AS-PATHs manually, and we find 7 AS-PATHs are suspicious, from which 3 AS-PATHs indeed have invalid ASNs, and 4 AS-PATHs have very short-lived AS links (cannot seen in later's CAIDA AS relationship data). It is noteworthy that the low suspicions are almost false alarms. In summary, if we consider all AS-PATH in RIBs to be real, then the accuracy for the detection of these AS-PATHs reaches 99.5% (6972/7000).

2) Misconfiguration and BGP Poisoning: Type-1 misconfiguration and BGP poisoning are very easy to detect, but Type-2 misconfiguration only get accuracy of 77.5%. This is because Type-2 misconfiguration generates many registered but unused ASNs which bypass Type-1 rule.

3) Path hijacking: In general, the results of our evaluation of Metis show that it performs well in detecting Type-1 hijacking with a satisfactory accuracy of 85.0%. However, its performance decreases for longer fake AS-PATH segments, as seen in Type-2 hijacking (accuracy of 50.7%) and Type-3 hijacking (accuracy of 30.0%). This is because as the length of the fake AS-PATH segment increases, the first AS in the fake segment is closer to the Internet core and is more likely to be a large ISP that is interconnected with other ASes, making the fake AS link more reasonable. For example, in the Type-2 hijacking case "59919 6939 32505 53561, 174

TABLE I
RESULT OF CRAFTED AS-PATHS

Type of AS-PATH	Number	Reliable link	Type-1 link	Type-2 link	valid AS-PATH	Suspicious AS-PATH				Accuracy
						Type-1	high	medium	low	
GREEN AS-PATHs	7000	11181	358	187	6966	5	3	6	**20**	**99.5%**
Type-1 Misconfiguration	1000	2231	108	985	167	0	924	0	0	92.4%
Type-2 Misconfiguration	1000	2174	496	582	256	247	528	0	0	77.5%
Type-1 hijacking	1000	2213	163	940	125	3	345	481	46	**87.5%**
Type-2 hijacking	1000	3018	153	984	493	2	322	176	7	50.7 %
Type-3 hijacking	1000	3706	160	935	700	0	250	50	0	30.0%
Type-1 BGP poisoning	1000	2237	236	940	107	14	879	0	0	89.3%
Type-2 BGP poisoning	1000	2241	372	2731	11	15	974	0	0	98.9%

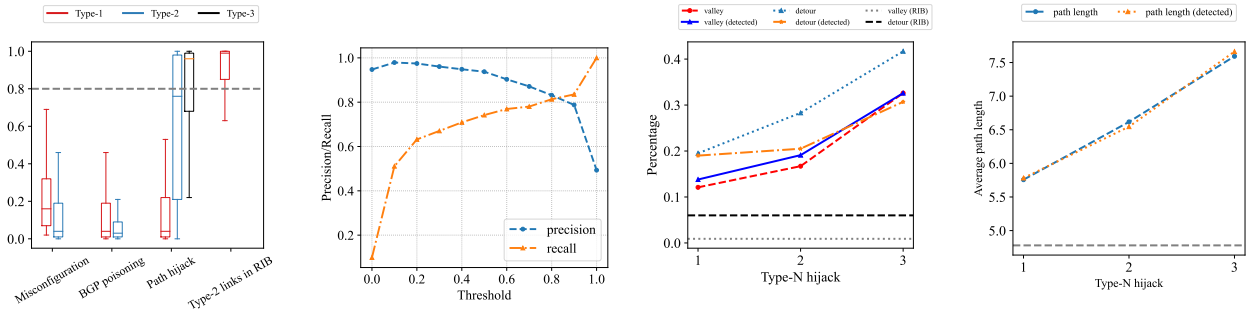

(a) (b) (c) (d)

Fig. 6. (a) Distribution of prediction value of Type-2 links. (b) Precision and Recall curve. (c) Detour and vally Path ratio of Type-N hijacking (including All AS-PATHs and the AS-PATHs detected by Metis). (d) Path length of Type-N hijacking

395554" (the comma separates the real and fake segments), AS53561 (Packet Forensics, US) is an attacker attempting to hijack AS39554 (Fullrate, Denmark). It inserts "174 395554" in the end and then announces the illegitimate route to its neighbor AS32505 (Conterra, US). As a result, an unseen AS link "53561 174" is created. Our link predictor output 0.99 for it, i.e., the model considers this link very likely to be real. We know that 174 is Cogent Communications, whose AS rank is 3. There may exist a link between 53561 (Packet Forensics, US) and 174 (Cogent Communications, US).

However, we argue that attackers will rarely use path hijacking with a large N. Because as "N" grows, the hijacking will make it more obvious due to the valley path and detour path. In Figure 6(c), the percentage of the AS-PATHs that have hegemony valley and country detour increase rapidly with increasing N. In all Type-3 hijackings, 32.6% AS-PATHs have a valley, and 41.7% have a country detour. And for comparison, the ratio of AS-PATHs that have valley and detour in RIB of 2021-11-01 0:00:00 are only 0.9% and 6%, respectively. Furthermore, path hijacking with large N will decrease the propagation range due to its low-attractive longer AS-PATH (Figure 6(d) show the average AS-PATH length increases proportionally with N) and the inserted high-tier AS not accepting the fake route (AS-PATH loop self-check mechanism).

D. Metis vs Argus

Metis can be fully integrated into any unseen link based path hijacking detection system, including Argus (the state-of-the-art). To illustrate how much overhead we can save

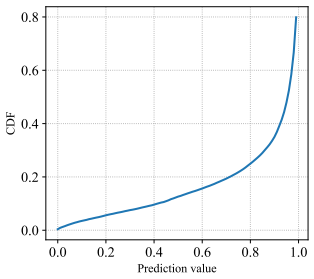

Fig. 7. CDF of prediction value of Type-2 links in November 2021.

if we integrate Metis into Argus. We run Metis with BGP updates of RRC00 throughout November 2021. To eliminate the effects of a cold start, so we let Metis run for one day and then calculate the unseen links from 2nd, November. On average, we can receive 161,808.2 reliable links, 30 new AS, 244.0 Type-1 links (7.3 suspicious Type-1 links), and 1,321.0 type-2 links per day. Figure 7 show the CDF of the prediction value of all Type-2 links. Suppose we set the prediction threshold t as 0.8. Then the ratio of suspicious Type-2 links is 0.229. Totally, Metis filter 236.7 Type-1 links and 1,018.5 Type-2 links. If each unseen link corresponds to one probe, Metis can save 80.2% cost. And if we only validate highly suspicious AS-PATHs, it will save even more.

V. RELATED WORK

Link Prediction. Link prediction is an old topic that is particularly well-researched. But there are few studies on

link prediction on AS topology. Recently, Zhuang et al. [24] define the prediction of unseen links of AS topology as a matrix-completion problem and reach a maximum AUC of 0.834. Kirtus G. Leyba et al. [25]. take into account the errors in the BGP routing data and use a statistical inference method for inferring the AS network topology. Their works focus on inferring complete AS topology using link prediction. While our work is not to infer the AS topology but to evaluate the authenticity of unseen links with link prediction.

Path hijacking Detection.

Existing state-of-the-art detection algorithms on path hijacking are based on unseen links. Argus [3] takes the AS-PATH of links that have not been seen in the past two months as anomalies and then verifies them through the data plane. Similarly, [12] assume if an AS link has never been observed in previous route announcements or a few prefixes use routes traversing this edge, it is highly suspicious. Artemis [11] directly treats AS links that have not appeared in the last ten months as anomalies. Metis is also an unseen link based method but can filter a lot of normal unseen links.

VI. Conclusion and Future work

To improve on current methods for detecting path hijacking, we propose an approach that models the evaluation of unseen links as a link prediction problem on AS topology. We then implement a real-time fake AS-PATHs detection system Metis. Our future work includes exploring more reliable AS library construction methods and link prediction algorithms and adding data plane probing to verify potential path hijacking events.

Acknowledgment

This research is supported by the National Key Research and Development Program of China, 2020YFE0200500. Congcong Miao is the corresponding author.

References

[1] Y. Rekhter, T. Li, and S. Hares, "Rfc 4271: A border gateway protocol 4 (bgp-4)," 2006.

[2] X. Zhao, D. Pei, 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, 2001, p. 31–35.

[3] X. Shi, Y. Xiang, Z. Wang, X. Yin, and J. Wu, "Detecting prefix hijackings in the internet with argus," in *Proceedings of the 2012 Internet Measurement Conference*, 2012, pp. 15–28.

[4] "bgpmon," 2021. [Online]. Available: https://bgpstream.com/

[5] "Bgpwatch," 2021. [Online]. Available: https://bgpwatch.cgtf.net/

[6] "Youtube hijacking: A ripe ncc ris case study," 2008. [Online]. Available: https://www.ripe.net/publications/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study

[7] "Chinese isp hijacks the internet," 2010. [Online]. Available: https://www.bgpmon.net/chinese-isp-hijacked-10-of-the-internet/

[8] H. Ballani, P. Francis, and X. Zhang, "A study of prefix hijacking and interception in the internet," *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 4, pp. 265–276, 2007.

[9] "How 3 hours of inaction from amazon cost cryptocurrency holders $235,000," 2021. [Online]. Available: https://arstechnica.com/information-technology/2022/09/how-3-hours-of-inaction-from-amazon-cost-cryptocurrency-holders-235000/

[10] L. Qin, D. Li, R. Li, and K. Wang, "Themis: Accelerating the detection of route origin hijacking by distinguishing legitimate and illegitimate MOAS," in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 4509–4524.

[11] P. Sermpezis, V. Kotronis, P. Gigis, X. Dimitropoulos, D. Cicalese, A. King, and A. Dainotti, "Artemis: Neutralizing bgp hijacking within a minute," *IEEE/ACM Trans. Netw.*, vol. 26, no. 6, p. 2471–2486, dec 2018.

[12] X. Hu and Z. M. Mao, "Accurate real-time identification of ip prefix hijacking," in *2007 IEEE Symposium on Security and Privacy (SP'07)*. IEEE, 2007, pp. 3–17.

[13] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang, "Phas: A prefix hijack alert system." in *USENIX Security symposium*, vol. 1, no. 2, 2006, p. 3.

[14] "Bgp hijacks observatory," 2021. [Online]. Available: https://dev.hicube.caida.org/feeds/hijacks/

[15] "The caida as relationships dataset, <1998-2022>," 2021. [Online]. Available: http://www.caida.org/data/active/as-relationships/

[16] J. M. Smith, K. Birkeland, T. McDaniel, and M. Schuchard, "Withdrawing the bgp re-routing curtain: Understanding the security impact of bgp poisoning via real-world measurements," *arXiv preprint arXiv:1811.03716*, 2018.

[17] L. Lü and T. Zhou, "Link prediction in complex networks: A survey," *Physica A: statistical mechanics and its applications*, vol. 390, no. 6, pp. 1150–1170, 2011.

[18] C. Orsini, A. King, D. Giordano, V. Giotsas, and A. Dainotti, "Bgpstream: a software framework for live and historical bgp data analysis," in *Proceedings of the 2016 Internet Measurement Conference*, 2016, pp. 429–444.

[19] M. Zhang and Y. Chen, "Link prediction based on graph neural networks," *Advances in neural information processing systems*, vol. 31, 2018.

[20] M. Zhang, Z. Cui, M. Neumann, and Y. Chen, "An end-to-end deep learning architecture for graph classification," in *Thirty-second AAAI conference on artificial intelligence*, 2018.

[21] "Special-purpose as numbers," 2021. [Online]. Available: https://www.iana.org/assignments/iana-as-numbers-special-registry/iana-as-numbers-special-registry.xhtml

[22] L. Gao, "On inferring autonomous system relationships in the internet," *IEEE/ACM Transactions on networking*, vol. 9, no. 6, pp. 733–745, 2001.

[23] S. Cho, R. Fontugne, K. Cho, A. Dainotti, and P. Gill, "Bgp hijacking classification," in *2019 Network Traffic Measurement and Analysis Conference (TMA)*. IEEE, 2019, pp. 25–32.

[24] S. Zhuang, J. H. Wang, J. Wang, C. An, Y. Xu, and T. Wu, "Predicting unseen links using learning-based matrix completion," in *NOMS 2022-2022 IEEE/IFIP Network Operations and Management Symposium*, 2022, pp. 1–9.

[25] K. G. Leyba, J. J. Daymude, J.-G. Young, M. E. J. Newman, J. Rexford, and S. Forrest, "Cutting through the noise to infer autonomous system topology," in *IEEE INFOCOM 2022 - IEEE Conference on Computer Communications*, 2022, pp. 1609–1618.