



Review

The state of affairs in BGP security: A survey of attacks and defenses

Asya Mitseva*, Andriy Panchenko, Thomas Engel

University of Luxembourg, 6 Avenue de la Fonte, Esch-sur-Alzette L-4364, Luxembourg



ARTICLE INFO

Keywords:

Border Gateway Protocol (BGP)
Internet routing
Security

ABSTRACT

The Border Gateway Protocol (BGP) is the de facto standard interdomain routing protocol. Despite its critical role on the Internet, it does not provide any security guarantees. In response to this, a large amount of research has proposed a wide variety BGP security extensions and detection-recovery systems in recent decades. Nevertheless, BGP remains vulnerable to many types of attack. In this work, we conduct an up-to-date review of fundamental BGP threats and present a methodology for evaluation of existing BGP security proposals. Based on this, we introduce a comprehensive and up-to-date survey of proposals intended to make BGP secure and methods for detection and mitigation of routing instabilities. Last but not least, we identify gaps in research, and pinpoint open issues and unsolved challenges.

1. Introduction

The Internet consists of a large set of networks, called autonomous systems (ASs), identified by unique AS numbers (ASNs). Each AS contains a number of hosts and routers under the administrative control of a single entity [1]. ASs are granted a set of delegated IP addresses, which are, in turn, assigned to their hosts and routers. Each AS is responsible for forwarding traffic at least to and from its IP addresses. To do this, ASs are connected via dedicated links and negotiate reachability information using Border Gateway Protocol (BGP). Contiguous IP addresses are aggregated in blocks (i.e., *prefixes*), each consisting of a n -bit IP address (n equals 32 for IPv4 and 128 for IPv6 respectively) and a mask length to reduce the number of routes distributed between ASs [1].

ASs set up routes via BGP based on their local policies. Despite its critical role on the Internet, BGP is not designed to provide any security guarantees and, thus, remains vulnerable to attacks [2,3] and misconfigurations [4], causing instabilities in the routing system or severe reachability problems. As well as frequent routing incidents [5,6], BGP threats are exploited for country-level censorship [7,8], damage to cryptocurrencies [9], and tracking users of anonymization networks [10]. Moreover, spammers often conduct BGP attacks so as to remain untraceable [11,12].

Although a large amount of research has proposed a variety of BGP security extensions in recent decades [13–15], none has been universally deployed on the Internet. Besides the lack of AS willingness, most of the solutions either only target particular aspects of routing vulnerabilities or require a significant computational overhead [16,17].

In this work, we conduct an up-to-date review of fundamental BGP vulnerabilities and possible attacks. We also bring together various requirements for BGP security solutions based on previous research and present a methodology for evaluation of existing security proposals. Last but not least, we introduce an extensive and up-to-date overview of proposals intended to make the routing protocol secure.

Early surveys [18–21] provided a review of BGP security solutions focusing only on some of the issues. Later works [22,23] presented a more broad overview. In 2010, Butler et al. [24] published the most comprehensive survey of both research and standardization efforts to make BGP secure. Concurrent work was performed by Huston et al. [25]. In contrast to our work, these works covered BGP security extensions reflecting the state of the art more than ten years ago and only a limited number of detection-recovery approaches. Because of the critical role of BGP on the Internet and its inadequate protection against attacks, research on securing BGP continued. Thus, a new comprehensive and up-to-date survey is needed to systematize and reveal the potential and the limits of the state-of-the-art solutions.

Recently, Siddiqui et al. [26] reviewed the standardization efforts made by the IETF Secure Inter-Domain Routing Working Group (SIDR WG) to secure BGP, but focused on two major security extensions only. Cardona et al. [27] focused on contributions intended to improve only the availability of BGP networks. Al-Musawi et al. [28] reviewed various techniques seeking to detect BGP anomalies, but neither presented methods for localization of root causes and routing anomaly mitigation nor focused on any security solutions. The most recent survey by Benesby da Silva [29] summarized state-of-the-art efforts improving the interdomain performance without discussing BGP security issues. In

* Corresponding author.

E-mail address: asya.mitseva@uni.lu (A. Mitseva).

contrast, we present a broad overview of countermeasures to protect BGP against attacks proposed in recent years and survey methods for detection, mitigation, and localization of routing instabilities. We further summarize research exploring different properties of existing proposals. This group of works was typically overlooked by previous surveys [22,24–26].

Our **contributions** are as follows: (i) We review fundamental BGP threats and possible attack vectors (Section 3). (ii) We collect desired properties for secure BGP and present a methodology for evaluation of existing methods to securing BGP (Section 4). (iii) Based on this, we present an extended review of existing BGP security proposals and detection-recovery systems, identify their strengths and limits (Section 5 and 6), and highlight open issues, gaps in research and unsolved challenges (Section 7). In the following section, we introduce fundamentals of the Internet and BGP needed for the rest of this paper.

2. Basics of Border Gateway Protocol

Currently, the AS interconnectivity is mainly based on confidential business agreements. In recent decades, the inferring of relations between ASs has been extensively studied by the researchers [30–33]. According to Gao et al. [34], there are two main groups of commercial relationships: customer-provider and peer-to-peer. In a customer-provider relationship, an AS (i.e., customer) pays another AS (i.e., provider) in order to access the rest of the Internet. The provider, in turn, may also be a customer of another AS. In a peer-to-peer relationship, two ASs with networks of similar size negotiate a link between themselves and exchange traffic free of charge. Based on the type of relationship with its neighbor, an AS decides how to propagate incoming routes to other ASs. The following export rules¹ are typically used: (i) Routes learned by a customer AS are advertised to other customer, peer, and provider ASs, and (ii) routes learned by a peer or a provider AS are announced to customer ASs only. Due to the greater resulting revenue, an AS always prefers a route via its customer over a route via its provider or peer.

On the Internet, most of the customer ASs only carry traffic whose source or destination is located within the AS, and are known as *stub* ASs. A stub AS that is connected to a single AS is called *single-homed* stub AS. However, stub ASs often establish connections to multiple ASs to provide resilience and load balancing for their services. Such ASs are known as *multi-homed* stub ASs. Provider ASs are also able to forward traffic whose source and destination are not in these ASs. These are called *transit* ASs.

Based on their routing policies and negotiated interconnections, ASs announce routes to a variety of prefixes through BGP [1]. Initially, routers (i.e., *BGP speakers*) establish a TCP connection from one to another and send an *OPEN* message to start a BGP session. To keep the session active, *KEEPALIVE* messages are periodically transmitted. Once the initial routing data has been exchanged between the routers, only incremental updates are sent when a routing information base (RIB) changes. To announce new routes, route updates, or route withdrawals, routers exchange *UPDATE* messages. In addition, the routers may employ *routing flap damping* (RFD) and *minimum route advertisement interval* (MRAI) timers. The RFD timer measures how often a route is withdrawn and re-announced. If a given threshold is exceeded, the route is flagged as *damped* (i.e., unstable) and cannot be selected as the best route. The MRAI timer defines the amount of time that needs to pass before a route can be re-advertised to neighbors. The aim of both timers is to decrease the frequency of routing changes and message overhead and, thus, reduce the routers' load and limit overall routing instability.

Each *UPDATE* message contains the following main attributes: (i) A prefix advertised, (ii) an AS originating this prefix, (iii) a sequence of

ASs that the message traverses, called the *AS path*, and (iv) the IP address of the next router to the destination. In case of a multi-homed AS, the AS may include a *multi-exit discriminator* (MED) to specify the best link to its neighbors. When forwarding *UPDATE* messages, each router prepends the ASN belonging to its AS to the AS path. Thus, routing loops can be detected and eliminated.

On receipt of a valid *UPDATE* message, a BGP speaker stores the new routing information in its RIB. The RIB of every router consists of three distinct parts: Adj-RIBs-In, Loc-RIB, and Adj-RIBs-Out. The Adj-RIBs-In contains routes learned from other BGP speakers. In accordance with its local policy, the BGP speaker selects a set of routes that it will use itself, and stores them in Loc-RIB. The BGP router also chooses a set of routes that will be advertised to neighboring BGP speakers, and lists them in Adj-RIBs-Out. Although each AS defines its own routing policy, a general rule in the route decision process is that higher preference is given to a route with a more specific IP prefix, known as *longest prefix match rule* [1].

3. Classification of attacks on BGP

BGP threats are caused by three fundamental vulnerabilities. First, BGP infrastructure is susceptible to physical attacks by outsiders, e.g., damage to hardware or cables between ASs [3]. Such attacks fall into the field of physical and logical security and are out of scope of this work.

Second, neither BGP nor the underlying protocols include any mechanism that prevents tampering with protocol data by outsiders. Since BGP messages are carried upon a TCP session, approaches proposed to secure TCP connections (e.g., by using cryptography) can also be applied in the context of BGP. Due to space constraints, we do not discuss TCP attacks and possible countermeasures, but refer the reader to [35–43].

Third, even if an intentional corruption of control messages by outsiders can be eliminated by hardening the TCP protocol and physical links, BGP does not ensure that legitimate participants do not use protocol data in a malicious manner or distribute bogus data injected into the routing information. E.g., bogus attributes indicating a false origin AS or a tampered AS path may cause severe disturbance to the routing process. Protection against tampering with routing information is described as *securing the control plane*. Furthermore, BGP does not guarantee that routers forward packets in a manner consistent with the announcements they have made via control messages, i.e., the packets can be dropped, rerouted or delayed. Therefore, *securing the data plane* is also required.

Focusing on the third group of threats, we summarize different attacks based on previous research and present the following attack taxonomy:

Data falsification attacks: A malicious AS is capable of injecting corrupted routing data into BGP messages. Here, the following attack vectors are possible:

Prefix hijack: As shown in Fig. 1(a), an AS falsely claims to originate a prefix not delegated to it. This causes a multiple origin AS (MOAS) conflict to be observed by other ASs. E.g., on November 6, 2015, a large Indian Internet service provider (ISP) started originating thousands of foreign prefixes.² The bogus announcements were accepted by some ASs and further propagated to their neighbors. Although MOAS does not directly indicate an attack [44], the attacker can avoid conflict by originating unadvertised prefix (e.g., used by spammers). A recent study [12] has shown that more than 20% of the global prefix space is delegated, but not publicly announced.

Subprefix hijack: Another way the attacker can prevent a MOAS conflict is by advertising a subnetwork of an existing prefix which does not belong to the attacker (Fig. 1(b)). This event is also known as de-

¹ Also known as *valley-free* rules [34].

² <https://bgpmon.net/large-scale-bgp-hijack-out-of-india/>.

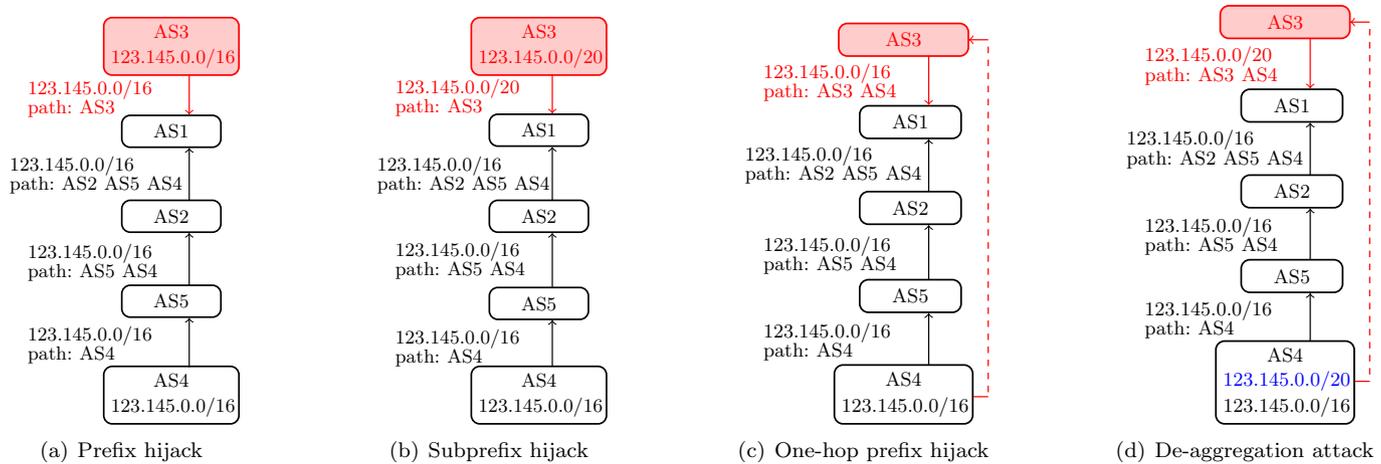


Fig. 1. Types of (sub)prefix hijack attacks: The malicious AS3 falsely announces the (sub)prefix of AS4 and attempts to force AS1 to adopt fake routes to this (sub)prefix.

aggregation attack.³ If no other ASs originate this prefix, most ASs adopt the route due to the longest prefix match rule. E.g., on April 27, 2017, a large Russian ISP started claiming to originate several more specific prefixes of existing prefixes that are typically advertised by other ASs.⁴

AS path forgery: The attacker may arbitrarily tamper with the AS path in UPDATE messages. Instead of forging the origin AS, he modifies the AS path to avoid a MOAS conflict and causes *one-hop prefix hijack*. To do this, the attacker announces a fake link between his AS and the victim AS (Fig. 1(c)). Another version of this attack is to announce a fake link to a subprefix of the victim AS (Fig. 1(d)), also known as *one-hop subprefix hijack*. Research by [47,48] showed the practicability of these attacks. In addition, due to economic incentives ASs may also intentionally modify the AS path in BGP messages and, thus, advertise more attractive (e.g., shorter) routes at the control plane, but still use another sequence of ASs at the data plane to forward the traffic. This attack is known as *traffic attraction attack* and its feasibility was empirically explored in [49].

Interception attack: This is an improved version of (one-hop) (sub)prefix hijacks. The attacker has a valid route to the victim AS. He can not only redirect traffic through it, but also forward it back to the real destination without disturbing the connectivity. On December 12, 2017, for a time period of six minutes a Russian AS conducted several prefix and subprefix hijacks of IP blocks belonging to well known and high traffic Internet organizations and rerouted the attracted traffic back to the legitimate destinations.⁵

Replay/Suppression attack: A malicious AS replays or suppresses withdrawal for a previously announced route. Although there is no such real event documented, every registered Internet outage may be caused by this attack.

Collusion attack: Two colluding non-neighbor ASs create a virtual tunnel between each other and build a BGP session through it. Thus, they generate forged routes without causing any suspicious routing conflicts. The feasibility of the attack was empirically demonstrated in [50].

Protocol manipulation attacks: Here, a malicious AS aims to manipulate properties of the routing protocol itself. The following attack vectors are possible:

MED modification: Similar to other BGP attributes, the multi-exit discriminator is not also protected which, in turn, may lead to tampering with MED values of routes. Thus, a malicious AS may affect ASs'

³ Please note that ASs often apply de-aggregation to recover their prefixes from prefix hijack [45] or implement traffic engineering and, thus, steer traffic through specific paths [46].

⁴ <https://bgpmon.net/bgpstream-and-the-curious-case-of-as12389/>.

⁵ <https://bgpmon.net/popular-destinations-rerouted-to-russia/>.

decisions.

Exploit RFD/MRAI timer: A malicious AS artificially withdraws and re-announces a route. Thus, ASs using the RFD timer consider the route unstable and ban it. On the other hand, ASs employing the MRAI timer delay their distribution of the corresponding UPDATE messages. Hence, this route may seem unreachable for some ASs. The feasibility of the attack was empirically showed in [51].

Data misuse attacks: An AS uses correct routing data in a malicious way. Possible attack vectors are as follows:

Denial of service (DoS): By causing heavy congestion on routers or links carrying BGP messages [2,52], the attacker creates *congestion-induced BGP session failures*. When BGP sessions are recovered, the routers first need to exchange full routing tables. This increases their load [53] and introduces significant convergence⁶ delays [2]. Another type of DoS is to create continuous withdrawals and re-advertisements of target routes to a victim AS, causing *deliberate link flapping*.⁷ Other ASs tag these routes as unstable and start suppressing their further propagation. Although there are no such real events documented, every registered Internet reachability problem may be caused by these attacks.

Route leak: This attack occurs when an AS propagates routes to ASs not intended to receive them under the terms of negotiated business agreements. E.g., a customer AS leaks a route received from one provider to another, even though this contradicts with the valley-free export rules (cf. Section 2). On August 26, 2017, Google accidentally leaked routes learned from its peers to some of its providers and, thus, became a transit AS [54]. As a result, many users (especially those in Japan) experienced slowness of the Internet or even a complete connectivity disruption.

To sum up, attacks on BGP increase the load on routers and cause instabilities and connectivity problems [2,4,55]. E.g., the adversary may create a *back hole* (i.e., drop traffic) by hijacking a prefix or eavesdrop traffic and perform man-in-the-middle attacks in case of an interception.

4. Desired properties for secure BGP

Previous research took into account distinct requirements that need to be addressed when developing a BGP security solution. To better understand the strengths and limits of existing BGP security proposals, we bring together all desired properties based on previous research that

⁶ BGP convergence measures the propagation and processing time needed by routers to settle on a best route.

⁷ In contrast to the protocol manipulation attacks, this attack does not require the RFD timer to be in use.

Properties	Approach																
	Route filtering	IRR	Use of DNS	IRV	RPKI	RPKI enhanced	S-BGP	soBGP	psBGP	BGPsec	S-BGP enhanced	SPV	SMPC-based routing	Blockchain-based routing	Listen & Whisper	Use of traceroute	
Security																	
Control-/ Data-plane Attacks Covered:	■ / □	■ / □	■ / □	■ / □	■ / □	■ / □	■ / □	■ / □	■ / □	■ / □	■ / □	■ / □	■ / □	■ / □	■ / □	■ / □	■ / □
Prefix / Subprefix hijack	□ / □	■ / ■	■ / ■	■ / ■	■ / ■	■ / ■	■ / ■	■ / ■	■ / ■	■ / ■	■ / ■	■ / ■	■ / ■	■ / ■	■ / ■	■ / ■	□ / □
AS path forgery	□	■	□	■	□	■	□	■	□	■	□	■	□	■	□	□	□
Interception attack	□	■	□	■	□	■	□	■	□	■	□	■	□	■	□	□	□
Replay/Suppression attack	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□
Collusion attack	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□
MED modification	□	□	□	■	□	□	■	■	□	□	□	□	■	□	□	□	□
Exploit RFD/MRAI timer	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□
Denial of service	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□
Route leak	□	■	□	■	□	□	□	□	□	□	□	□	■	■	□	□	□
Performance																	
Convergence delay	□	□	□	■	□	□	■	■	■	■	■	□	■	■	□	■	■
Stability	■	■	■	□	■	■	□	□	□	□	□	□	□	□	□	□	□
Scalability	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□
Computational overhead	□	□	□	□	□	□	■	■	■	■	■	■	■	■	■	■	■
Bandwidth overhead	□	□	□	■	□	□	■	■	■	■	■	■	■	■	■	■	■
Storage overhead	■	■	■	■	■	■	■	■	■	■	■	■	□	■	■	■	■
Privacy																	
Routing privacy	□	□	□	□	□	□	□	□	□	□	□	□	■	■	□	□	□
Deployability																	
Deployability	■	■	■	□	■	■	■	■	■	■	■	■	□	□	■	■	■
Adoptability	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□
Status																	
Adopted / Standardized	□ / □	□ / □	□ / □	□ / □	□ / □	□ / □	□ / □	□ / □	□ / □	□ / □	□ / □	□ / □	□ / □	□ / □	□ / □	□ / □	□ / □
Academic paper only	□	□	■	■	□	■	■	■	■	■	■	■	■	■	■	■	□

Fig. 2. Summary of main BGP security proposals. RPKI enhanced represents first-hop authorization approaches and S-BGP enhanced describes the use of amortization and aggregation of signatures. Notion: ■ has feature, □ partially has feature, □ does not have feature.

need to be incorporated in the new secure routing protocol. Our categorization covers four main dimensions: security, privacy, performance, and deployability.

4.1. Security properties

BGP security proposals need to take into consideration all known vulnerabilities of the legacy protocol (described in Section 3) and provide countermeasures to eliminate these threats. Moreover, they need to guarantee that they do not introduce any new attack vectors that have not been observed when using the plain BGP protocol.

4.2. Performance properties

BGP security solutions need to provide certain performance properties to be deployed and used on the Internet.

Convergence delay: How fast do routers settle on a best route? As a baseline, we assume the convergence time measured when the legacy protocol is used.

Stability: Besides the risk of a prolonged convergence time, it is important to verify if convergence is ever reached. Otherwise, the quality of connections to some prefixes may degrade or even be lost during the transient period.

Scalability: How well does a new secure routing protocol scale as the number of ASs adopting it increases? Does the protocol scale when fully deployed on the Internet?

Computational overhead: How many BGP control messages need to be processed per unit of time? How many additional, CPU-intensive, operations need to be executed on each message compared to ordinary BGP? This also considers the case when auxiliary, more powerful hardware is needed to perform these computations.

Bandwidth overhead: This refers to the rate of BGP control messages to be sent. Furthermore, including additional attributes in the messages to provide security impacts the bandwidth required to transmit them.

Storage overhead: This refers to the degree of additional memory on each BGP router that a new BGP security solution requires compared

to the plain BGP implementation. It also includes the case when an AS needs to use auxiliary hardware to fulfill the new storage requirements.

4.3. Privacy properties

BGP security proposals should also consider the fact that ASs often desire to keep their routing policies, business relations, and other commercial data private.

Routing privacy: How much additional sensitive information does the new secure protocol reveal compared to the legacy protocol?

4.4. Deployability properties

The integration of a new BGP security solution depends not only on its degree of security and performance properties, but also on how easy it is to bootstrap the solution on the Internet scale. To explore this, we consider the following properties.

Deployability: Can the new BGP security proposal be deployed incrementally over a period of time? Does it enable information forwarding between routers supporting the secure protocol and those running the legacy protocol in the same AS? Due to the large number of ASs composing the Internet, the coexistence of a plain and a secure BGP implementation should be taken into account. This means that ASs that have upgraded their infrastructure to the new protocol also need to provide backward compatibility to enable routing with non-adopters.

Adoptability: Once the incremental adoption of a new BGP security solution is initiated, adoptability denotes the quantity of volunteer ASs willing to adopt the new protocol over time, i.e., how far the adoption will spread. The set of initial adopters and their routing policies significantly impact the degree of protocol adoptability.

5. BGP Security enhancements

As early as 1988, Perlman [56] conducted the first comprehensive study on routing protocol security. She drew attention to the fact that routing protocols are robust to simple failures, but cannot counteract Byzantine failures, i.e., if a node intentionally modifies, delays or forges

messages. Previous surveys [22,24,25] extensively reviewed early research works suggesting improvements for BGP and other distance-vector routing protocols. Nevertheless, since these proposals neither presented a complete BGP security solution nor took into account deployment and performance issues, finding defenses to protect BGP against attacks remained an active area of research. In the following, we review current state-of-the-art BGP security solutions and discuss their pros and cons with respect to the set of properties presented in Section 4. Fig. 2 summarizes main BGP security solutions presented in this section. While we still describe some early approaches also covered by [22,24,25], we use them as a baseline to show the reader how each group of BGP security solutions has evolved in recent decades. Recently, Siddiqui et al. [26] and Al-Musawi et al. [28] also presented surveys on BGP security. Contrary to our work, they either reviewed the standardization efforts made by the IETF SIDR WG only or focused on detection techniques and completely ignored BGP security proposals. Last but not least, we summarize works that mainly concentrated on the evaluation of different properties of existing proposals. This area of research was typically overlooked by previous surveys [22,24–26].

5.1. Real-world protection practices

Before focusing on research proposing BGP security enhancements, we present current protection practices used by network operators. Today, several ASs extend their routing policies by adding rules to filter out potentially bogus prefix announcements. Network operators take into account not only AS business relationships, but also filter out special-use IP addresses, announcements containing private ASNs or too long AS paths [24]. They often restrict advertisements for networks smaller than /24 to prevent size explosion of the global routing tables [57] and the number of prefixes that a neighboring AS can announce [58]. Thus, the routers are protected from memory exhaustion and a possible leaking of entire network tables is limited. To create their route filters, network operators also rely on publicly available lists of bogus announcements [59] and public repositories with routing data. Route filtering is particularly effective if applied by provider ASs to restrict bad routes announced by their customers. Recently, network operators have additionally applied peer locking [60] to further restrict faulty routes unintentionally announced between provider ASs and peer ASs. To construct AS path filters, they rely on out-of-band communication with each other to announce allowed interconnections.

Due to their simplicity and effectiveness, route filtering and peer locking are an attractive incomplete solution used by many ASs for many years. ASs neither need to modify the legacy protocol nor invest money in hardware replacement. However, the creation, maintaining, and update of filter lists creates a significant computational and storage overhead. Apart from this, the need for exchange of commercial data seriously violates the privacy requirement and may discourage many ASs from using peer locking. While peer locking limits possible misconfigurations, it cannot prevent intentional attacks as it is based on trust between ASs. Thus, neither approach provides a long-term solution for secure interdomain routing.

A popular public repository containing routing data is the *Internet Routing Registry* (IRR).⁸ IRR aims to provide a shared global view of correct routing information by enabling ASs to voluntarily upload routing data. Other ASs can retrieve this data and use it to build their route filters. Despite tools developed to execute consistency checks of registered policies [61,62], uploading and extracting data from IRR remain error-prone due to complex routing policies and the distributed nature of the IRR [4,63]. Moreover, the IRR itself has to be secure, accurate, and up-to-date to be effective. Due to lack of access level privileges, ASs are not willing to upload confidential routing information. Although previous research [64] tried to enhance the IRR by

enabling authorized verifiable post and search of routing data, many security concerns regarding content integrity and authorization needed to make changes to the IRR remain unsolved.

5.2. Proposals securing the control plane

The primary goal of most of the works is to ensure that the data payload of the BGP protocol and the semantics of the payload are authentic and correct. In other words, a router can reliably verify that incoming data is not removed, modified or replayed during the transmission. On a receipt of an UPDATE message, the router has to be able to validate the advertised prefix, authorize the origin AS, and check the correctness of the AS path and other additional attributes if present. Existing proposals use various techniques (including cryptographic) to do this.

5.2.1. Proposals relying on DNS

Reference to [65] suggested the use of the DNS as a distributed database for origin validation. To this end, the authors assumed a creation of a new DNS zone for IP prefix delegation managed by IANA. Within this zone, each node in the DNS tree corresponds to one or more IP prefixes delegated to an AS. On receipt of an UPDATE message, a BGP router performs a DNS query to verify the origin data announced in the message. However, this approach does not replicate the IP address allocation hierarchy and, thus, introduces management issues. In response to this, recently Gersch and Massey [66,67] presented a new naming convention where the mask length of the prefix is encoding in a binary format. Although these works assume the use of DNSSEC to provide data authenticity, the main issue with this idea is that it tries to solve BGP security problems by placing implicit trust in another insecure protocol. Moreover, these proposals do not address route authorization.

5.2.2. Proposals relying on overlay networks

Several works have suggested keeping BGP unchanged, but use complementary protocols or technologies to validate control-plane routing data. Early work by Goodell et al. [68] proposed *Interdomain Route Validation* (IRV) protocol that operates independently of BGP. Every AS manages its own IRV server, where it stores its AS-specific policy data. On receipt of an UPDATE message, a BGP router contacts its local IRV server to check if the obtained routing data is correct. The IRV server, in turn, queries the IRV server of the relevant AS to validate the routes. Hence, to verify the correctness of an AS path, one queries the IRV servers of all ASs involved. Each IRV server can enforce levels of access control over sensitive routing data and, thus, limit the exposure of confidential relationships. BGP speakers can also upload routing reports to IRV servers, including received announcements, topology data, etc. Thus, the IRV service is able to detect misconfigurations and monitor network health in general. Recently, Chen and Haeberlen [69] enhanced the IRV protocol to provide a secure communication between IRV servers. To this end, the authors introduced differentially private query-answer sessions. Nevertheless, the main limitations of these approaches remain their scalability, the need for AS collaboration, bootstrapping the service, and recovering from outages.

To reduce the complexity of BGP and, thus, the quantity of misconfigurations, an amount of research went one step further and suggested to completely separate the routing decision process from the routers [70]. Kotronis et al. [71,72] outsourced the routing control plane to an external trusted entity, called *contractor*. As several ASs will select the same contractor, each contractor can build clusters of ASs and maintain them in a centralized manner using a Software Defined Networking (SDN) controller. Other works [73,74] designed SDN-based Internet Exchange Points (IXPs) to improve the operation of the Internet routing, e.g., by implementing more expressive traffic control policies, without removing BGP. While these works mainly aimed to improve the Internet performance, they did not explore the impact on the routing

⁸ <http://irr.net/index.html>.

security.

To sum up, solutions taking use of supplementary protocols or technologies neither require modifications to BGP nor create further computational and storage overhead for, or require reconfiguration of, BGP speakers. However, they are less dynamic than routing changes. This may create problems when route authorizations change quickly and slow down convergence. ASs also need to set up additional equipment to create an overlay network.

5.2.3. Proposals relying on asymmetric cryptography

The most promising group of BGP security proposals is based on asymmetric (i.e., public-key) cryptography. Typically, they rely on Public Key Infrastructure (PKI) for assignment and distribution of public keys.

Kent et al. [14] proposed the first comprehensive approach to secure BGP, *Secure-BGP* (S-BGP). S-BGP employs digital signatures and associated public key certificates to validate announced routing data. By using a PKI that imitates the existing IP address and ASN assignment system rooted by IANA, the authors authenticate prefix allocations, routers' identities, and organizations owning ASs. Each organization that has an allocated set of prefixes or ASNs has been issued a certificate by the resource holder in the PKI hierarchy. The resource holder also issues certificates for the routers belonging to the ASs. To verify that an AS is allowed to originate a route to an IP prefix, Kent et al. used digitally-signed statements, *address attestations* (AAs). An AA consists of a single AS and a set of IP prefixes, and is signed by the resource holder. The AAs are distributed out-of-band and verified through the certificate chain from the origin AS to IANA. AAs eliminate (sub)prefix hijack attacks, but do not prevent AS path modifications. Thus, *route attestations* (RAs) are added to the syntax of UPDATE messages. A RA is signed by each router which it was forwarded through, where all routers on the path sign previously attached signatures. Although RAs mainly target protecting the AS path attribute, they can also be extended to cover other vulnerable BGP attributes, e.g., the MED [75]. To secure the underlying TCP session and prevent active wiretap attacks from outsiders, the IPsec protocol [40] is used.

Although S-BGP addresses most of the BGP threats, it introduces significant computational costs and performance issues that hamper its adoption [76,77]. In response to this, White suggested *Secure Origin BGP* (soBGP) [15] aiming to provide a trade-off between security and overhead. Like S-BGP, soBGP uses a PKI consisting of three types of certificates. The first certificate type binds an ASN to a public key. Based on a web-of-trust authentication model, each AS has a certificate issued by a trusted anchor (e.g., large ISPs or well-known authentication service providers). The second certificate type binds an ASN to a set of prefixes that this AS is allowed to advertise. The third certificate type contains information about the routing policies and neighboring ASs of each AS. This data is used by a soBGP router to create its own view of the network topology, which, in turn, is needed to validate incoming UPDATE messages. If the AS path contained in the UPDATE message violates the router's topology, the route is dropped. Compared to S-BGP, where AAs are distributed out-of-band, in soBGP the certificates providing origin authentication are exchanged in-band via a new message type. The other two certificate types are distributed out-of-band. Thus, in contrast to S-BGP, where the RAs are sent with every UPDATE message, the network topology built by soBGP is static and requires frequent synchronization of topology updates across ASs. Although soBGP provides more configuration flexibility [78], the variety of options and the use of non-standard certificates create several interoperability challenges [75].

van Oorschot [79] and Wan et al. [80] presented *pretty secure BGP* (psBGP), which implements origin authorization based on a decentralized trust model between ASs, and validates the AS path in route announcements using a rating-based approach with the help of a certificate hierarchy. In psBGP, every AS creates a *prefix assertions list* (PAL) used to validate whether an AS is allowed to originate a prefix. A

PAL contains the set of prefixes delegated to that AS and a list of IP prefix ownership assertions for the neighbors from the point of view of that AS. To verify the authenticity of an origin AS, a BGP router checks the consistency between the PALs of ASs that are neighbors of the origin AS. Like S-BGP, psBGP uses signatures to perform AS path validation. Each AS also expresses its believe in the trustworthiness of the other ASs. To reduce the computational overhead, psBGP uses these beliefs to decide whether to validate all signatures on a path or only a subset of them.

Besides the significant protocol overhead created by S-BGP, soBGP, and psBGP, a major deployment barrier for these solutions is the non-existence of a global PKI. Hu et al. [81] suggested a strategy to incrementally develop such a PKI. First, each AS creates and uses a self-signed certificate to sign prefix announcements. As adoption increases, larger ASs sign the certificates of their customers and become trust anchors for them. Under the control of IANA, large ASs can cooperate to establish a well-defined root of trust. Also, there is the need to negotiate signing algorithms in router environments [22]. In 2012, the SIDR WG presented the first standardized architecture, *Resource Public Key Infrastructure* (RPKI) [82,83], which provides a global PKI for origin authorization in a similar manner to S-BGP.

While RPKI limits (sub)prefix hijack attacks, it does not prevent AS path modifications. *BGPsec* [84] is a proposal to ensure the authenticity and correctness of the sequence of ASs which a route announcement traverses. BGPsec relies on RPKI to ensure origin authorization. As with S-BGP, each router uses its certificate to sign the data received from the previous router, the ASN it belongs to, and the ASN of the next router on the path, and includes it in the message. BGPsec also requires periodic key rollovers of routers' certificates to restrict replay attacks by authorized routers [85]. This means routers' certificates are periodically refreshed and UPDATE messages are re-sent using the newly-generated certificates. As most of the ASs (i.e., stub ASs constituting about 85% of the ASs on the Internet [86]) need to replace their hardware in order to support BGPsec, the SIDR WG also allowed the use of an asymmetric BGPsec communication, *simplex BGPsec* [16]. A stub AS sends BGPsec announcements and receives plain UPDATE messages from its provider. Stub ASs do not need to perform heavyweight BGPsec signature validation and, thus, are encouraged to adopt BGPsec. Recently, the SIDR WG further improved BGPsec by inserting a special field in UPDATE messages that sets AS routing policy to prevent route leaks [87]. This field is set to zero to indicate no restrictions concerning the recipients of a route, and one to specify that providers and peers should not receive the route. Yet, it does not ensure that the policy is indeed enforced. The information about restrictions may also reveal commercial data. Song et al. [51] further proposed an improvement to BGPsec in order to limit protocol manipulation attacks, in particular the manipulation of RFD and MRAI timers. The authors added *secure root cause information* to each UPDATE message that is generated by the root cause AS and contains its ASN, timestamp, and prefix destination. To limit the frequency of intentional routing changes and, thus, keep overall routing stability, each receiving router verifies the announcement and discards all messages whose checks fail.

The benefits provided by BGPsec in partial deployment are questionable, which slows down its adoption [17]. To fill the gap meanwhile, Cohen et al. [48,88] proposed an extension to RPKI. Each resource holder issues a path-end record containing a list of adjacent ASs through which its AS is reached. On receipt of an UPDATE message, a BGP router uses these records to check if the penultimate AS appears in the list of neighbors of the origin AS. While this prevents one-hop prefix hijack attacks, two- and more-hop prefix hijack is possible. Although the authors argue that the success rate of such attacks decreases due to longer paths announced by the attacker, related work has shown that, even by advertising a longer path, an adversary is able to attract a significant amount of traffic [89].

The BGP security proposals described above are typically divided into four groups based on the degree of security they provide. Methods

that only validate if an AS is allowed to advertise a prefix, e.g., RPKI, are known as *origin authorization* (OA). *First-hop authorization* (OA + 1) represents approaches validating not only the prefix allocation, but also the first AS on the route adjacent to the origin AS, e.g., RPKI including path-end records. Both groups neither require protocol modifications nor create convergence delays in the Internet. However, they assume a supplementary hardware to validate prefix origin ASs and only protect against a few specific routing attacks. Along with origin authorization, methods that further verify the existence of a route to a prefix based on an AS-level Internet map, e.g., soBGP, are known as *routing topology path verification* (RTPV). The final group is called *path validation* (PV). Here, the methods provide not only origin authentication, but also validate every AS on the route to the announced prefix, e.g., S-BGP, psBGP, BGPsec. These are believed to be the most promising solutions, due to the extensive security they offer. Although intuitive optimizations to these proposals have been discussed, e.g., using additional hardware [76] or caching UPDATE messages to reduce the number of signature verification operations [76,77,90,91], their expensive cryptographic operations remain a major obstacle to their adoption.

In response to this, several works have focused on decreasing the protocol overhead created by these solutions. These works rely mainly on different cryptographic schemes to reduce the computational costs, either regarding prefix ownership authorizations or route attestations. Concentrating on the costs of origin authentication, Aiello et al. [92] and McDaniel [93] formalized the semantics of prefix delegation and proposed efficient proof structures for carrying delegation attestations. Besides simple attestations created for each route announcement (as in S-BGP) and authenticated delegation lists containing all delegations made by a single institution, the authors also considered authenticated delegation trees based on a Merkle hash tree and authenticated delegation dictionaries. Using the latter, they showed the feasibility of in-band origin authentication.

Nicol et al. [77] applied signature amortization to reduce the number of signatures needed. Each BGP speaker builds a bit vector representing intended receivers when announcing the same route to multiple neighbors. Instead of creating distinct signatures, the router signs the UPDATE message once, appends the bit vector, and sends it to its neighbors. Each receiving router uses the vector to check if it is the target recipient. To advertise several routes to different ASs, a BGP speaker first collects UPDATE messages intended for a set of ASs and creates a Merkle hash tree with them. The router signs them collectively by executing only one signing operation. Again, it sends the signature, the signing material needed for verification, and the bit vector in a single UPDATE message. By using the vector, a receiving router checks if the message is intended for it and validates the route. While this method is efficient in terms of processing time, it dramatically increases the size of UPDATE messages compared to S-BGP and, thus, the bandwidth overhead and router's memory costs.

Subsequent work by Zhao et al. [94] combined signature amortization with another cryptographic technique for aggregating signatures to reduce the memory overhead. The main advantage of the aggregate signature approach is that the amortized signature for a set of messages has the same length as a regular signature for one message. Recently, Brogle et al. [95] suggested another aggregate signature method that allows a BGP speaker to insert its own signature to an unverified UPDATE message, sends it immediately, and, thus, postpone the verification for some later time period. However, since each router has to also add a random string to the signature to enable the later verification, the signature size of an UPDATE message grows linearly with the number of ASs in the AS path.

To avoid the use of PKI, Boldyreva et al. [96,97] applied identity-based (IB)⁹ sequential aggregate signatures. However, IB-based systems

lack a method of key revocation [13]. In response to this, Mancini et al. [98] enhanced the approach by adding timestamps to the signatures and, thus, enforced prefix and route revocations. Another work by Li et al. [99] also suggested a method based on IB-based cryptography. Since the authors rely on a trusted attestation service running on each BGP router, they argue that each AS needs to only validate the announcements of its neighbors in order to guarantee AS-path authenticity. The main issue with this idea is that the routing security strongly depends on the security of the attestation service. Xiang et al. [100] used another method to reduce the computational costs for signing and verifying signatures. As ASs decide which routes to distribute based on business contracts with their direct neighbors, they need to sign and announce only path segments representing a route learned from a previous AS and exported to a successor AS.

To reduce the costs of path validation, Butler et al. [101] combined several methods, including hash chains and signature aggregation. A router sends all available paths to its neighbors along with tokens representing hash chain anchors. When a route changes, it can be represented by an authentication token that is verified through hash operations. Like [94], the authors used a Merkle hash tree to sign a set of UPDATE messages collected until the MRAI timer expires. Finally, they exploited the stability of path advertisements through caching cryptographic proofs and, thus, reduced the number of cryptographic operations needed.

In summary, this group of studies does not provide a complete BGP security solution, but tries to optimize distinct building blocks of existing proposals. While most of the works focus mainly on minimizing the computational and storage overhead, they inadequately evaluate the impact of the proposed optimization on the other properties, e.g., convergence delay, scalability, stability. Most of the works are based on a misinterpretation of the semantics of the MRAI timer [25], incorrectly assuming that this timer effects all UPDATE messages to a given router.

5.2.4. Proposals relying on symmetric cryptography

Another way to keep low protocol overhead when securing BGP is to apply symmetric cryptography. Research by Hu et al. [102] suggested a method for path validation using nested message authentication codes (MACs). Each AS listed in the AS path of an UPDATE message shares a secret key with a predefined node intended to verify the message (the *validator*). An origin AS builds a MAC over a concatenation of an initial authenticator value and its prefix, and adds the MAC to an UPDATE message. Analogously, each subsequent AS creates a new MAC, in which it uses the MAC obtained from the incoming announcement. Each following MAC covers not only the data received, but also the authenticator value of the previous router. On receipt of an UPDATE message, the validator uses all known secret keys to recursively verify the AS path announced.

This approach was improved in a follow-up work by Hu et al. [13]. The authors presented *Secure Path Vector (SPV)* protocol using a sequence of one-time off-line signatures, where the signer executes expensive cryptographic operations prior to their use, making the signing procedure faster. Although SPV is efficient in terms of computational costs, it creates significant overhead, as a large amount of state information needs to be transferred and processed. Moreover, Raghavan et al. [103] showed that SPV does not prevent AS path forgery and collusion attacks. The authors also argued that some standard digital signatures are as efficient as one-time signatures with respect to route validation. Based on this, Yin et al. [104] further optimized the performance of the signature generation and verification algorithms used in SPV, and also reduced the size of the created signatures.

Bruhadeshwar et al. [105] explored another symmetric key approach, based on square grid protocol, to overcome the complexity of SPV. An interesting aspect of this work is that the authors assume that one of the routers on the path is trusted in order to reduce the size of the advertisements. Once the trusted node has validated an UPDATE message, it does not need to forward the signing material from the previous BGP

⁹ IB cryptography is a type of public-key cryptography where the public keys of the parties are their identities. The corresponding private keys are maintained by a trusted third-party entity [98].

routers to its successors.

Nevertheless, the main concerns that prevent the use of symmetric key approaches to secure BGP remain the long-term security of the symmetric keys (i.e., they cannot be used for long periods as they are vulnerable to brute-force attacks), and their resistance to collusion.

5.2.5. Privacy-preserving techniques

Recent works have applied novel paradigms to secure the routing protocol and keep the AS routing policies private. Gupta et al. [106] proposed a method based on secure multi-party computation (SMPC) that outsources the route computation. To this end, the work assumes a predefined set of computational servers. Each AS sends its routing policy (i.e., route ranking and export policy) to the servers. The routing policy is divided into n parts (n equals the number of servers) and each part is sent to one server. The servers apply SMPC over this input to compute possible routes and send them back to those ASs involved in each route. While this proposal enables fast convergence and easy upgrades of the routing protocol, its scalability and computational overhead are questionable. In particular, the approach requires high processing time, as the internal computations executed by the servers mirror the BGP operation, i.e., the servers simulate BGP by using the input data from the ASs.

Asharov et al. [107] further explored the idea of using SMPC to secure interdomain routing. By applying an improved state-of-the-art SMPC algorithm, the approach ensures security in case of no honest majority and scales better for a large number of ASs. Compared to [106], each AS only sends information about its relations with neighbors. ASs may additionally label a preferred neighbor or define export policies to express special preferences. To reduce the computational costs, the algorithm does not consider stub ASs, as they do not have customers and do not carry transit traffic.

Henecka and Roughan [108] proposed another privacy-preserving routing protocol that computes the shortest AS path based on Bellman-Ford algorithm. Instead of outsourcing route computations, each AS forwards encrypted messages along all known links to a predefined destination, cumulatively summing weights for the links by using homomorphic encryption. On receipt of all messages, the destination AS can determine the shortest path, even though the intermediate ASs on the route are kept private. The main disadvantage of this work is its significant message overhead. It further requires a prior knowledge of candidate destinations for the creation of routes between ASs.

Inspired by the idea of Bitcoin, Hari and Lakshman [109] used the blockchain to secure interdomain routing. The authors introduce two types of transaction published in the blockchain. The first aims to provide prefix and ASN ownership authentication and authorization. If an address holder has been allocated a prefix, an attempt to allocate this prefix to another entity by anyone other than the address holder will fail. The second transaction type records all route announcements to provide path validation. The input of a current transaction corresponds to the previous AS on the path and the output represents the next AS of a route. However, the use of the blockchain introduces scalability issues due to the huge number of transactions, and slows down convergence as the transactions are published periodically on the blockchain.

To sum up, to make their proposals attractive for deployment, these works focused mainly on the privacy guarantees that a new secure routing protocol needs to offer. However, since these solutions are too young, issues such as computational overhead, scalability, and single point of failure due to the use of centralized servers, are not extensively explored and merit further research.

5.3. Proposals securing the data plane

The design of BGP does not include any mechanism to guarantee that routing announcements validated in the control plane match the actual forwarding paths in the data plane [110]. An intuitive solution to this issue is to check path consistency between the control plane and the

data plane through *traceroute* when forwarding traffic. By sending packets with increasing time to live (TTL) values and receiving the corresponding ICMP time exceeded responses from a destination, *traceroute* creates a sequence of addresses on the path as long as no packets are being lost. However, besides the lack of message integrity in *traceroute* (i.e., an attacker may tamper with probe packets), Mao et al. [111] discovered difficulties in retrieving AS-level paths from *traceroute* output. To mitigate this, the authors proposed an AS-level *traceroute* tool that relies on a comprehensive set of IP-to-AS mappings created by multiple publicly available sources, e.g., *whois*, BGP routing data from specific vantage points. Padmanabhan and Simon [112] further enhanced *traceroute* by adding cryptography to ensure the integrity and authenticity of data-plane probes. Another work by Augustin et al. [113] suggested *Paris traceroute*, which aims to mitigate non-existent nodes and loops falsely reported by *traceroute*. To sum up, *traceroute*-based countermeasures are appropriate for diagnostics when a problem is detected, but they do not scale for path consistency checks in real time, as they need to permanently probe all possible routes.

Instead of using *traceroute*, Avramopoulos and Rexford [114] proposed encrypted tunnels between routers, through which data traffic and probes are forwarded. While this method verifies end-to-end connectivity, it does not provide full route integrity, i.e., traffic still can be misrouted by intermediate ASs. Another work [115] presented a combination of techniques to address this. The authors assume that the intermediate ASs are explicitly specified by the source AS in every data packet. A MAC authentication mechanism is used for packet validation. In addition, the destination AS needs to acknowledge the receipt of every packet to the source and intermediate ASs. In summary, apart from the fact that both proposals require AS collaboration and prior distribution of shared secrets (or keys), they also create significant computational and bandwidth overhead and, thus, may not scale if deployed on the Internet.

To reduce this overhead but still keep a certain degree of route integrity, Wong et al. [116] proposed a dedicated lightweight verification protocol. An AS initiating a transmission injects a part of a predefined shared secret into some of the data packets. An AS receiving this data traffic replies to the sending AS by using the other part of the shared secret to prove its presence on the path. The routers also keep a list of valid and invalid responses to detect and localize potential malicious BGP speakers. While this proposal avoids the use of cryptographic operations, it relies on expensive off-line secret exchange. To enable on-line distribution of secrets, Bruhadeshwar et al. [117] presented key distribution protocols for [116]. Similar source and path validation methods were also presented in [118–122]. Typically, these works enforce the routers to add a cryptographic primitive to each packet and, thus, ensure that the traffic traverses the correct path. For example, Liu et al. [119,120] aimed to provide source authentication by inserting tagged pieces of information into IP packets. A source AS adds a sequence of ASNs and the corresponding MACs generated using predefined secret keys shared between the source AS and the other ASs on the route. Each succeeding AS validates its MAC using its shared key. Nevertheless, the primary limitation of these approaches remains scalability, the need for prior distribution of shared secrets (or keys), and computational and bandwidth overhead. In addition, they do not address incremental or partial adoption.

Research by Subramanian et al. [123] suggested the *Listen & Whisper* protocols that detect not only inconsistencies in the data plane (*Listen*), but also address control-plane vulnerabilities (*Whisper*). The *Whisper* protocol relies on nested signatures added to route announcements to provide origin authorization. Each origin AS creates a signature field in an UPDATE message to bind a particular prefix to its ASN. Each intermediate AS updates this field by using a cryptographic hash function. The *Listen* protocol inspects TCP packets to detect data-plane anomalies. In particular, it observes the completeness of TCP connections, i.e., if the transmission of management packets is followed by data packets, and compares data packets and acknowledgments to identify a possible

Research work	Target solution	Properties			
		Security	Performance	Privacy	Deployability
Kent et al. [76]	S-BGP	–	✓	–	–
Nicol et al. [77]	S-BGP, signature amortization	–	✓	–	–
Chan et al. [132]	OA protocols, OA+1 protocols, RTPV & PV protocols	–	–	–	✓
Avramopoulos et al. [133]	S-BGP, soBGP	–	–	–	✓
Raghavan et al. [103]	SPV	✓	–	–	–
Goldberg et al. [110]	PV protocols	✓	–	–	–
Goldberg et al. [89, 144]	OA protocols, soBGP, S-BGP, data-plane verification, route filtering	✓	–	–	–
Gill et al. [16]	S-BGP, soBGP, BGPsec	–	–	–	✓
Boldyreva et al. [145]	S-BGP, soBGP	✓	–	–	–
Lychev et al. [17]	RPKI, S-BGP, BGPsec	–	–	–	✓
Cooper et al. [139]	RPKI	✓	–	–	–
Heilman et al. [140]	RPKI	✓	–	–	–
Wählisch et al. [137]	RPKI	–	–	–	✓
Gilad et al. [138]	RPKI	–	–	–	✓

Fig. 3. Overview of research works investigating different properties of a target BGP security solution.

prefix reachability problem. Mizrak et al. [124,125] and Argyraki et al. [126] also monitored traffic patterns to detect faulty packet forwarding. Based on traffic statistics that are periodically collected and shared between ASs, network operators can localize the source of packet loss, delay or anomalous traffic. Again, these methods require significant storage and processing overhead, and introduce scalability restrictions. Moreover, they do not use any cryptographic operations, making them vulnerable to bogus traffic injection and tampering. Other works focused on generalized data-plane fault localization protocols, which are also adaptable to BGP. Due to space constraints, we do not discuss them, but refer the reader to [127–129].

Zhu et al. [130] presented an alternate design for interdomain routing. Instead of utilizing BGP, routers create an approximate topology map of AS interconnectivity by using active and passive probing in the data plane. Wendlandt et al. [131] suggested building a route repository, used to determine an optimal path to a destination. The authors do not directly assume the truthfulness of the routes. Instead, the service measures the performance of each path (e.g., by using *Listen* [123]) to identify if it is a usable route. Large transit ASs may additionally provide information about existing routes to the repository. As in [114], the service creates a tunnel to the destination domain to transmit data traffic.

In conclusion, all data-plane methods described above (expect *Listen & Whisper*) are auxiliary countermeasures and need to be used in combination with a control-plane proposal. The existing data-plane methods also rely mainly on an inspection of distinct packets by verifying cryptographic primitives, making them unsuitable for widespread deployment.

5.4. Evaluation of existing proposals

Instead of suggesting new BGP security enhancements, several works concentrated on investigating different properties of existing proposals. Early work by Kent et al. [76] evaluated the performance of S-BGP and empirically showed the significant computational, bandwidth, and storage overhead of the protocol. Nicol et al. [77] explored the convergence delay introduced by pure S-BGP and S-BGP using a signature amortization technique. Beside the performance issues, there is no consensus on how to bootstrap more sophisticated secure routing protocol on Internet scale. Therefore, several works evaluated the degree of security, deployability, and adoptability of distinct proposals. Fig. 3 summarizes of this amount of research.

Chan et al. [132] were the first who showed the importance of the adoptability in the design of BGP security solutions when partially deployed. Given a predefined set of initial adopters, they showed that PV methods indicate up to ten times greater adoptability under weak attacker model¹⁰ compared to OA and RTPV. OA+1 and PV provide similar degree of adoptability under strong attacker model¹¹. Yet, these results are biased by the simplified AS topology model used for simulations, i.e., ASs always prefer shortest paths and announce their paths to all neighbors. Furthermore, Chan et al. did not discuss how to select the initial set of adopters. To overcome the latter, Avramopoulos et al. [133] studied the security impact of different groups of ASs originally adopting a security solution. They achieved a good protection of traffic transmitted by the adopters using up to five Tier-1 ASs.¹² The authors also argued that control- or data-plane detection methods can be added to the security protocol to get around the problem of partially-secure paths. Gill et al. [16] further proposed a strategy how to efficiently adopt security solutions, e.g., S-BGP, soBGP or BGPsec, by creating proper incentives for ASs to deploy them. Like [133], Gill et al. confirmed the need for a few Tier-1 ASs initially adopting the new protocol. Although these works [16,132,133] showed promising results regarding partial deployment, they assume that ASs always prefer secure routes.

Contrary to the assumption above, Gill et al. [135] revealed that the majority of network operators do not prioritize the use of secure routes in their routing policies. Taking this into account, Lychev et al. [17] showed that a secure BGP solution such as S-BGP and BGPsec has only meager benefits in partial deployment in contrast to the standardized solution, RPKI. Even worse, the need for a coexistence of plain and secure BGP implementation creates new attack vectors. On the one hand, malicious ASs may intentionally disable the use of the secure BGP for some routes, known as *downgrade attack*, and, thus, render the deployment of secure BGP useless for groups of ASs. On the other hand, the lack of consensus amongst ASs where to place the security in the routing policy may cause the existence of multiple stable BGP states.¹³

¹⁰ The attacker can only access BGP traffic sent directly to it.

¹¹ The attacker can eavesdrops any BGP connection.

¹² Tier-1 ASs compose the core of the Internet. They do not have providers and peer with all other Tier-1 ASs [134].

¹³ If the primary route chosen by an AS is broken, the AS selects another route. When the initial route is available again, the AS still keeps the route later selected instead of converging to the previous stable BGP state.

Contrary to [16,133], Lychev et al. showed that Tier-2 ASs¹⁴ initially adopting the secure protocol will provide better security than Tier-1 ASs.

The only one standardized security solution, RPKI, still suffers from poor adoption on the Internet despite the extensive efforts devoted by the SIDR WG and the support by major router vendors [136]. Several works explored potential reasons for this scarce adoption. Wählisch et al. [137] showed that only 6% of the prefixes containing Alexa¹⁵ Top one million websites are covered by RPKI. The more popular websites are less likely to be secured than the less popular. The authors revealed that some organizations worry to deploy RPKI, as its certificate hierarchy indicates to some extent AS business relationships. Gilad et al. [138] argued that the main factors that hinder the RPKI adoption are human errors, causing the invalidity of RPKI objects, and dependencies between organizations, causing ASs wishing to incorporate RPKI to wait for other ASs to deploy it first. They showed that about 10% of the IP address space is upward-dependent, i.e., an AS incorporating RPKI should receive a RPKI object signed by an entity which did not deploy RPKI yet, and about 90% of the prefixes are downward-dependent, i.e., an AS incorporating RPKI should wait another entity allowed to advertise a subprefix of the AS's prefix to obtain a RPKI object first.

Along with the lack of incentives for adoption, the hierarchical architecture of RPKI assigns to centralized authorities disproportionate power to unilaterally revoke authorization or tear down prefixes under their control. In particular, Cooper et al. [139] and Heilman et al. [140] demonstrated how abusive organizations may cause legitimate routes to be classified as invalid by manipulating the contents of RPKI publication points and, thus, rejected by ASs preferring secure routes only. Last but not least, concerns regarding the RPKI scalability, e.g., total number of RPKI objects, consistent, fresh view of all RPKI objects by ASs, have been raised in recent years [141–143].

Several works have shown that promising BGP security proposals such as RPKI, S-BGP, and BGPsec partially solve the security problems of the legacy protocol even *fully* deployed. Goldberg et al. [89,144] explored the weaknesses of these secure protocols and quantified their efficacy against traffic attraction attacks. The authors showed that BGP security proposals only address the security of route announcement semantics, but cannot ensure the contractual legitimacy of the routes. In particular, Goldberg et al. argued that an adversary is able to attract even more traffic by only manipulating his export policies rather than using data falsification attacks. Another work by Goldberg et al. [49] further argued that unrealistically strong AS routing policy restrictions, along with a fully deployed control-plane security solution, e.g., S-BGP, are needed to ensure correct data-plane packet forwarding. Boldyreva and Lychev [145] performed a cryptographic analysis of secure routing protocols by studying, generalizing, and formalizing known threats. They formally proved that S-BGP satisfies origin and route authorization, but does not ensure route validity, and soBGP fails to meet both route validity and route authorization.

To sum up, while the majority of recent research mainly explore the deployability and security properties of several BGP security proposals, other features such as scalability and stability completely lack of investigation.

6. Detection and mitigation of anomalies

The fact that none of the existing BGP security solutions incorporates all desired properties (cf. Section 4), in particular the incomparability with existing hardware and the meager benefits achieved by state-of-the-art approaches when partially deployed, minders the adoption of a new secure routing protocol for future time. As a temporal

workaround, some research has focused on detecting, localizing, and mitigating routing anomalies. While previous surveys either reviewed BGP security solutions and ignored anomaly detection mechanisms [18,26] or covered a limited quantity of detection and mitigation techniques reflecting the state of the art more than ten years ago [22,24,25], we look through modern state-of-the-art methods aiming to not only identify and mitigate routing anomalies, but also locate the root cause. Recently, Al-Musawi et al. [28] reviewed different methods seeking to detect BGP anomalies. Contrary to our work, the authors neither surveyed works trying to locate the perpetrator of a suspicious routing event nor reviewed mitigation proposals. In the following, we briefly summarize early research¹⁶ on detection, localization, and mitigation techniques and focus on more recent works not covered by previous surveys in detail. Thus, we aim to show how the research in this area has evolved in the recent decade.

Anomaly detection: Detection techniques aim to discover suspicious information or behavior in routing data and then raise alarms. We divide existing detection methods into three groups based on the type of information they monitor: control-plane methods, data-plane methods, and a combination of both. The *control-plane detection* approaches usually observe UPDATE messages to recognize malicious intents. Early works mainly focused on detecting (sub)prefix hijack events. Some proposals either relied on a static prefix ownership map or used optional attributes in BGP messages to identify bogus announcements [24]. Research by Qiu et al. [146] suggested the origin AS to be directly notified by the ASs receiving multiple origins advertisements for its prefix. Lad et al. [147] presented a real-time alert system where the address owners register their prefixes. The system monitored announcements from public databases and notified prefix owners via email if AS origin data changes. Qiu et al. [53] revealed a similar method based on historical routing data of prefix ownerships and AS link data. Various heuristics were applied to validate the collected information. Siganos and Faloutsos [148] used prefix allocations and ASN assignments from regional Internet registries and AS routing policies from IRR to verify the correctness of origin data announced in UPDATE messages.

To avoid a single point of failure created by the systems described above, Hiran et al. [149,150] detected (sub)-prefix hijacks using a distributed method. The authors applied two Chord-based distributed hash tables (DHTs) containing prefix origin data and information related to the AS itself, respectively. Similar to [147], this method relied on reports of participating ASs to build the ground truth. On a receipt of an announcement containing a new (i.e., not seen before) prefix or origin AS, the receiving AS queries the first DHT to check for a potential (sub)prefix hijack. If the event is confirmed, the legitimate AS is notified. The second DHT is mainly applied for collaborative fault detection, e.g., identifying spamming. Like [149,150], Heaberlen et al. [151] also used existing trust and business relationships between ASs to detect routing anomalies. BGP routers within an AS record all UPDATE messages they sent or received in a log file. The file, along with a set of rules describing the AS routing policy, is revealed to the AS's neighbors to check how the rules are followed. To prevent log tampering, BGP is extended to support message authentication for logs and acknowledgements confirming the successful storing of the log by neighbors. Yet, this method requires ASs to reveal their routing policies. Similar methods for detecting AS relationship violations and anomaly traffic were also presented in [152–155]. Compared to [151], they rely on cryptographic signatures or SMPC, together with principal component analysis to avoid revealing commercial data.

Various machine learning methods were also used for routing anomaly detection. Research by Al-Rousan and Trajković [156] applied Support Vector Machines and Hidden Markov Model to identify malicious routing intents. The authors proposed several types of features to

¹⁴ In contrast to Tier-1 ASs, Tier-2 ASs need to purchase transit to reach some destinations on the Internet.

¹⁵ <https://www.alexa.com/>.

¹⁶ Interested readers are referred to [22,24,25,28] for more details.

represent a routing anomaly, e.g., the total number of announcements and withdrawals exchanged. Relying on features similar to those presented in [156], Guo et al. [157] applied an adaptive fusion-based method to build a model for normal routing behavior and used it for anomaly detection. Zhang et al. [158] relied on unsupervised clustering techniques to recognize deviations from the normal state of BGP data flow. Lutu et al. [159] argued that prefixes that are less visible in the Internet than expected by the prefix owners is an indicator for routing issues. The authors periodically collected BGP messages from vantage points and assigned a visibility degree to each prefix based on the fraction of ASs that have an active stable route for it. Given a ground-truth dataset representing the expected visibility status of prefixes provided by network operators, Lutu et al. applied a decision-tree machine learning model to automatically distinguish potential routing instabilities.

Recently, Li et al. [160] revealed the correlation between routing loops and route leaks and presented an approach for detecting routing loops and, thus, identifying active route leak events. Chang et al. [161] claimed that once an AS has conducted a prefix hijack, it is very likely that it will repeat this in some future time. By using knowledge about previous prefix hijack events, the authors generated reputation values for each AS and used them to trigger real-time alerts. Although network operators typically restrict the number of prefixes that a neighboring AS can announce, Khare et al. [58] empirically showed that routing-table-leak events occur in the Internet. To detect these events, the authors observed if an AS (e.g., that occasionally leaks its routing tables) starts simultaneously originating prefixes of many other ASs as its own for a short time period in order to detect potential prefix hijacks. Yet, this method generated a large number of false positives due to the similarity between prefix hijacks and legitimate routing changes. In response to this, Su et al. [162] focused on improving the detection scheme by Khare et al. The authors argued that typically prefix hijacks between neighboring ASs do not occur and Tier-1 ASs neither conduct nor experience prefix hijacks and, thus, significantly reduced the quantity of the false positives. Another amount of research [163–165] used visualization methods to graphically represent either UPDATE messages or changes of AS paths and prefix ownerships and, thus, drew attention on potential routing issues.

Data-plane detection methods typically involve permanent probing of the Internet to identify AS path modifications and prefix unreachability. Early work by Zheng et al. [166] used continuous active probing to measure the number of hops from a predefined vantage point (i.e., monitor) to a target prefix. If a significant variation is detected, the monitor further checks for changes in the AS path to a reference point. The reference point is an IP address outside of the target prefix, but still topologically very close to it. Thus, the route from the monitor to that reference point is almost identical (i.e., a sub-path) to the route to the target prefix. Hence, in contrast to valid route changes that affect both routes equally, prefix hijack introduces significant difference between them, as the IP address of the reference point is not within the affected prefix. Quan et al. [167] applied Bayesian inference to weigh periodically collected probing data and, thus, estimated the status of each prefix. If the availability of a prefix becomes uncertain, the system starts continuously scanning IP addresses belonging to this prefix until it determines whether the prefix is up or down. The use of Bayesian inference reduces the frequency of permanent active probing to target prefixes.

Contrary to the approaches above, Hiran et al. [168] used round trip time (RTT) measurements that are passively collected by end users in various networks while accessing Internet resources in order to avoid the overhead costs of active probing. The gathered data is periodically sent to a centralized entity that performs outlier detection to identify routing anomalies. Yet, major RTT deviations do not necessarily imply an attack. They can also occur due to network congestion. Therefore, beside RTT measurements Balu et al. [169] also took into account the number of intermediate hops, similarity of paths that packets traverse

on, and the propagation delay in order to reduce the quantity of false positives.

Zhang et al. [170] proposed a system where a prefix owner continuously probes known transit ASs to detect if its prefix has been hijacked. Djatmiko et al. [171] correlated network traffic from multiple monitors distributed in different ASs by using SMPC to prevent revelation of sensitive data. Each AS then individually analyzed the collected information to diagnose network outages. Goldberg et al. [172,173] focused on monitoring the performance of AS paths in the presence of an attacker. Instead of encrypting and authenticating all traffic [114], the sender and the receiver apply a pseudo-random hash function on each packet, store the output locally, but transmit packets unchanged. The sender and the receiver regularly exchange and compare the collected output to detect packet dropping activity. Yet, this approach does not verify AS identities on the path and cannot distinguish between malicious tampering and network congestion. Zhang et al. [174] focused on probabilistic acknowledgement-based protocols to localize a packet-dropping adversary while keeping a balance between improved detection rate and communication overhead. Like some data-plane detection methods [164,165], Fischer et al. [175] visualized collected probings of various destinations to identify routing anomalies.

Some related works combined *data- and control-plane measures* to overcome issues introduced by these methods when used separately. To detect (sub)prefix hijacks in real-time, early research by Hu et al. [47] applied data-plane probing to build fingerprints for end hosts and networks only after a suspicious UPDATE message is observed. To do this, Hu et al. considered several properties to characterize a prefix, e.g., operating systems of end hosts, TCP/ICMP timestamps. Research by Xiang et al. [176] and Shi et al. [177] used control-plane data to detect suspicious routes and data-plane measures to check the reachability of IP addresses within a prefix. If the system identified a potential malicious announcement, it used publicly-available looking glass servers¹⁷ to collect control-plane route status for the prefix and probed for active IP addresses belonging to this prefix. The authors compute the correlation between the control- and data-plane data gathered to recognize prefix hijack, i.e., most of the affected servers cannot receive a reply from the prefix, while most unaffected servers can. To detect interception attacks, Hiran et al. [149,150] compared the AS path announced in an UPDATE message with the AS sequence retrieved by actively probing the same prefix.

To prevent a large quantity of false positives, Schlamp et al. [178,179] tried to detect subprefix hijack attacks by using UPDATE messages to retrieve candidates of hijacked prefixes and conducting active scans of SSL/TLS-enabled hosts within these networks. If the hosts use different public keys before and during the potential subprefix hijack event, this event is considered as a real attack. Recently, Siddiqui et al. [180,181] presented a theoretical framework to model different types of route leaks and suggested methods how to detect each of them. Based on data from its routing table and knowledge about its business relationships with other ASs, an AS can identify potential route leaks from its customer ASs by checking whether incoming route announcements are valley-free. To recognize route leaks from its peer ASs, the victim AS combines active control-plane probes with passive data-plane traffic monitoring. The latter is also used to detect route leaks not directed at the victim AS, but still affecting its routes.

Currently, *BGPmon*¹⁸ is the most popular commercial detection tool widely used by many network administrators. By combining control- and data-plane data from a hundred vantage points worldwide, it identifies state changes of the Internet routing operation and monitors the reachability of prefixes. BGPmon triggers notifications to the network operators that are sent via multiple communication channels.

¹⁷ <http://www.bgp4.as/looking-glasses>.

¹⁸ <https://bgpmon.net/>.

However, the AS administrators should manually investigate each alert in order to determine whether it is a real routing instability.

Localization of anomaly source: Once a routing anomaly is detected, only a few works explored the problem of localizing the perpetrator of a routing failure. Early research by Feldmann et al. [182] argued that the local source of routing instability is located either on the path announced before or on the new path advertised during the routing event. By collating UPDATE messages announcing AS path changes from different vantage points, the authors created a set of suspicious ASs appearing in all routes that changed. Research by Javed et al. [183] analyzed path changes hop by hop and filtered those ASs as possible root causes whose next-hop downstream AS was different compared to the one announced before the alternation.

By using various heuristics, Katz-Bassett et al. [184] compared historical probe data with a current traceroute measuring to determine the potential misbehaving BGP speaker. As the Internet routing is asymmetric, the authors used spoofed packets to probe a path in one direction without needing to probe the other and, thus, identified the broken route direction. However, they were not able to determine the hops along the reverse path. To address this, Katz-Bassett et al. applied reverse traceroute in their follow-up work [185]. Qiu et al. [186] argued that an AS conducting (sub)prefix hijack or interception attack cannot manipulate the portion of a route between its upstream neighboring AS and a vantage point. Observing that the routes from a set of vantage points to an affected prefix converge around the misbehaving AS, the authors ranked candidate ASs based on their appearances in path neighborhood sets and the total distance between each of them and the vantage points in order to narrow down a set of potential misbehaving ASs.

Anomaly mitigation: Several works went one step further and explored possible mitigation techniques. Early research mainly focused on methods ignoring or demoting suspicious routes once they are detected. Wang et al. [187] suggested preferring known stable routes over short-lived routes to protect routes to top-level DNS servers. However, this method does not scale for arbitrary routes. Zhang et al. [188] contacted a set of predefined ASs, *lifesaver* ASs, that delete bogus routes from their routing tables and prevent their further propagation. Each lifesaver AS also announces shorter (more attractive) valid routes to the victim. Careful selection of lifesaver ASs is critical for the effectiveness of the bogus-route purging method and the promotion of valid routes. Karlin et al. [189,190] proposed *pretty good BGP* (pgBGP), which blocks the propagation of suspicious new paths for a given time period. While this prevents the distribution of bad short-lived routes, large convergence delays ensue. Therefore, Zhang et al. [191] suggested that ASs use trusted paths for traffic forwarding, but still propagate the suspicious routes to their neighbors.

Recently, Katz-Bassett et al. [185] exploited BGP poisoning to avoid traversing a misbehaving AS without disturbing other working routes. Given that an alternative path exists during a routing failure, the affected origin AS artificially creates a loop along the route containing the faulty AS. To avoid the loop, the misbehaving AS is implicitly forced to reject this route and send route withdrawals to its neighbors. Yet, this method is not applicable if the BGP loop prevention is disabled or ASes discard poisoned announcements. Holterbach et al. [192] suggested that BGP speakers redirect traffic through alternative paths when outages of primary routes are detected. To do this, BGP routers within a single AS attach a tag to each incoming packet containing the list of ASs which the packet will be forwarded through and a precomputed backup next-hop AS that is used in case of link failure. Sermpezis et al. [193] used prefix de-aggregation when possible or outsourced the distribution of recovery BGP announcements advertising the hijacked prefix to a third-party organization. By creating MOAS events, the organization attracts traffic addressed to the affected AS and tunnels it back this AS. Qiu et al. [194] discussed how to choose optimal locations for organization's agents to improve the gain from the mitigation method.

Summary: Various techniques for anomaly detection, localization,

and mitigation have been proposed to date. Nevertheless, none of them provides a complete detection-recovery system, i.e., they only tackle a specific subproblem. The methods are either impractical due to increased computational, bandwidth or storage overhead and need for inter-AS communication or rely on third-party data, which, in turn, make them inaccurate and easily compromised.

7. Discussion

As we have shown, BGP security solutions proposed to date either eliminate most of the BGP threats at the cost of high overhead, or sacrifice security goals to achieve performance. Even though methods based on signatures and attestations have received considerable attention, and some of these are partially standardized, frequent expensive cryptographic operations are a major obstacle to their adoption [76]. In response to this, some works attempted to reduce the number and length of signatures by exploiting specific protocol timers. However, these protocol timers are not used by the vast majority of router vendors and network operators at all [135], making these approaches infeasible. Instead of public-key cryptography, other works applied symmetric cryptography for signature generation. These proposals are still vulnerable to collusion attacks and do not reliably protect route announcements against AS path forgery attacks [103]. Recently, a few works have explored interdomain routing architectures based on SMPC. Nevertheless, high computational overhead and convergence delay, scalability, and single point of failure due to the use of central servers, remain issues and need further research.

Most of the BGP security solutions still focus on securing the control plane and cannot ensure that data traffic is indeed transmitted over negotiated routes. Unfortunately, the threat of subverting packet forwarding has received comparatively little attention. The existing data-plane solutions introduce a high protocol overhead, making them not applicable in reality. Network operators are unwilling to be transparent to probing tools such as traceroute, as their internal network structure can be revealed. Thus, they often block such unauthorized probing. If probing can be detected by a malicious AS, it can be treated differently from regular traffic to cover the attack. In addition, none of the existing BGP security proposals is able to counteract collusion attacks. Given a virtual tunnel between two non-neighboring ASs, they can distribute forged routes even if a secure routing protocol is fully deployed. Therefore, we encourage the community rethink the extensive effort devoted to adopt a new secure BGP protocol. Is it reasonable to force ASs to invest money in hardware replacement if the new protocol cannot solve almost none of the issues? We must also draw attention to the fact that most of the works proposing and evaluating BGP security solutions are based on empirically-driven AS topologies. Typically, these topologies are incomplete and not sufficiently accurate [195,196] which, may bias reported results.

Although the detection, mitigation, and localization of routing anomalies is an active field of research, none of the existing methods can identify and differentiate among different routing attacks in real-time. Moreover, most of these approaches rely on inaccurate and scarce datasets [195,197] and are prone to false alarms. Even if a malicious event is detected in a timely manner, the response of the ASs may take hours [148]. Thus, the damage increases, since the bogus routes are quickly distributed to other ASs.

Last but not least, recent studies [198,199] showed that BGP blackholing is often used for Distributed Denial of Service (DDoS) mitigation. By using BGP community attributes¹⁹ or relying on out-of-band information, ASs affected by DDoS attacks inform their neighbors the traffic from which prefixes needs to be dropped, i.e., which prefixes will be temporarily unreachable. Despite the popularity of this service,

¹⁹ The community attribute is an optional attribute of variable size that carries additional information in UPDATE messages [1].

none of the current state-of-the-art detection, localization, and mitigation methods take use of this information.

8. Conclusion

In this work, we have exposed fundamental drawbacks of the currently-existing interdomain routing protocol. In particular, we presented an up-to-date list of attacks on BGP, which pose a severe security risk for diverse applications and services on the Internet. This has allowed us to categorize and survey the protection and detection mechanisms for BGP proposed in the literature. Our analysis revealed that these approaches solve only a neglectable small fraction of the problems, and that in most cases at the cost of high overhead. Based on this, we finally discussed open issues and unsolved challenges for further research.

Acknowledgements

Parts of this work have been funded by the Luxembourg National Research Fund (grant no. C15/IS/10486741) (FNR) within the CORE Junior Track project PETIT.

References

- [1] Y. Rekhter, T. Li, S. Hares, A Border Gateway Protocol 4 (BGP-4), (<https://tools.ietf.org/html/rfc4271>).
- [2] O. Nordström, C. Dovrolis, Beware of BGP attacks, *ACM SIGCOMM Comput. Commun. Rev.* 34 (2) (2004) 1–8.
- [3] S. Frey, Y. Elkhatib, A. Rashid, K. Follis, J. Vidler, N. Race, C. Edwards, It Bends But Would It Break? Topological Analysis of BGP Infrastructures in Europe, European Symposium on Security and Privacy (EuroS&P), IEEE, Saarbrücken, Germany, 2016, pp. 423–438.
- [4] R. Mahajan, D. Wetherall, T. Anderson, Understanding BGP misconfiguration, Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, ACM, Pittsburgh, Pennsylvania, USA, 2002, pp. 3–16.
- [5] BGPStream, (<https://bgpstream.com/>).
- [6] A. Toonk, Large Hijack Affects Reachability of High Traffic Destinations, (<https://bgpmon.net/large-hijack-affects-reachability-of-high-traffic-destinations/>).
- [7] S. Khattak, T. Elahi, L. Simon, C.M. Swanson, S.J. Murdoch, I. Goldberg, SoK: making sense of censorship resistance systems, Privacy Enhancing Technologies (PETs), (2016), pp. 37–61. Darmstadt, Germany
- [8] M.C. Tschantz, S. Afroz, Anonymous, V. Paxson, SoK: towards grounding censorship circumvention in empiricism, Symposium on Security and Privacy (S&P), IEEE, San Jose, CA, USA, 2016, pp. 914–933.
- [9] M. Apostolaki, A. Zohar, L. Vanbever, Hijacking Bitcoin: routing attacks on cryptocurrencies, Symposium on Security and Privacy (S&P), IEEE, San Jose, CA, USA, 2017.
- [10] Y. Sun, A. Edmundson, L. Vanbever, O. Li, J. Rexford, M. Chiang, P. Mittal, RAPTOR: routing attacks on privacy in Tor, The 24th USENIX Conference on Security Symposium, USENIX Association, Washington, D.C., USA, 2015, pp. 271–286.
- [11] A. Ramachandran, N. Feamster, Understanding the network-level behavior of spammers, Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM), ACM, Pisa, Italy, 2006, pp. 291–302.
- [12] P.-A. Vervier, O. Thonnard, M. Dacier, Mind Your Blocks: on the stealthiness of malicious BGP hijacks, Symposium on Network and Distributed System Security (NDSS), (2015). San Diego, CA, USA
- [13] Y.-C. Hu, A. Perrig, M. Sirbu, SPV: secure path vector routing for securing BGP, Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, ACM, Portland, Oregon, USA, 2004, pp. 179–192.
- [14] S. Kent, C. Lynn, K. Seo, Secure border gateway protocol (S-BGP), *IEEE J. Sel. Areas Commun.* 18 (4) (2000) 582–592.
- [15] R. White, Securing BGP through secure origin BGP (soBGP), *Internet Protocol J.* 6 (3) (2003) 47–53.
- [16] P. Gill, M. Schapira, S. Goldberg, Let the market drive deployment: a strategy for transitioning to BGP security, SIGCOMM conference, ACM, Toronto, Ontario, Canada, 2011, pp. 14–25.
- [17] R. Lychev, S. Goldberg, M. Schapira, BGP security in partial deployment: is the juice worth the squeeze? SIGCOMM Conference, ACM, Hong Kong, China, 2013, pp. 171–182.
- [18] S. Goldberg, Why is it taking so long to secure internet routing? *Commun. ACM* 57 (10) (2014) 56–63.
- [19] S. Murphy, BGP Security Protections, 2002, (<https://tools.ietf.org/html/draft-murphy-bgp-protect-00>).
- [20] M. Yannuzzi, X. Masip-Bruin, O. Bonaventure, Open issues in interdomain routing: a survey, *IEEE Netw.* 19 (6) (2005) 49–56.
- [21] M. Zhao, S.W. Smith, D.M. Nicol, The performance impact of BGP security, *IEEE Netw.* 19 (6) (2005) 42–48.
- [22] M.O. Nicholes, B. Mukherjee, A survey of security technologies for the border gateway protocol (BGP), *IEEE Commun. Surv. Tutorials* 11 (1) (2009) 52–65.
- [23] T. Wan, P.C. van Oorschot, E. Kranakis, A selective introduction to border gateway protocol (BGP) security issues, NATO Advanced Studies Institute on Network Security and Intrusion Detection, IOS Press, 2005.
- [24] K. Butler, T.R. Farley, P. McDaniel, J. Rexford, A survey of BGP security issues and solutions, *Proc. IEEE* 98 (1) (2010) 100–122.
- [25] G. Huston, M. Rossi, G. Armitage, Securing BGP – a literature survey, *IEEE Commun. Surv. Tutorials* 13 (2) (2011) 199–222.
- [26] M.S. Siddiqui, D. Montero, R. Serral-Gració, X. Masip-Bruin, M. Yannuzzi, A survey on the recent efforts of the internet standardization body for securing inter-domain routing, *Comput. Netw.* 80 (2015) 1–26.
- [27] J.C. Cardona, P. Francois, B. Decraene, J. Scudder, A. Simpson, K. Patel, Bringing high availability to BGP: a survey, *Comput. Netw.* 91 (2015) 788–803.
- [28] B. Al-Musawi, P. Branch, G. Armitage, BGP Anomaly detection techniques: a survey, *IEEE Commun. Surv. Tutorials* 19 (1) (2016) 377–396.
- [29] R.B. da Silva, E.S. Mota, A survey on approaches to reduce BGP interdomain routing convergence delay on the internet, *IEEE Commun. Surv. Tutorials* 19 (4) (2017) 2949–2984.
- [30] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, kc claffy, G. Riley, AS Relationships: inference and validation, *ACM SIGCOMM Comput. Commun. Rev.* 37 (1) (2007) 29–40.
- [31] L. Gao, On inferring autonomous system relationships in the internet, Global Telecommunications Conference (GLOBECOM), IEEE, San Francisco, CA, USA, 2000, pp. 387–396.
- [32] V. Giotsas, M. Luckie, B. Huffaker, kc claffy, Inferring complex AS relationships, Conference on Internet Measurement Conference (IMC), ACM, Vancouver, BC, Canada, 2014, pp. 23–30.
- [33] A. Lodhi, N. Larson, A. Dhamdhere, C. Dovrolis, kc claffy, Using peeringDB to understand the peering ecosystem, *ACM SIGCOMM Comput. Commun. Rev.* 44 (2) (2014) 20–27.
- [34] L. Gao, J. Rexford, Stable internet routing without global coordination, *IEEE/ACM Trans. Netw.* 9 (6) (2001) 681–692.
- [35] S. Convery, D. Cook, M. Franz, An Attack Tree for the Border Gateway Protocol, (<https://tools.ietf.org/html/draft-ietf-rpsec-bgpattack-00>).
- [36] W.M. Eddy, TCP SYN Flooding Attacks and Common Mitigations, (<https://tools.ietf.org/html/rfc4987>).
- [37] V. Gill, J. Heasley, D. Meyer, P. Savola, C. Pignataro, The Generalized TTL Security Mechanism (GTSM), (<https://tools.ietf.org/html/rfc5082>).
- [38] M.G. Gouda, E.N. Elnozahy, C.-T. Huang, T.M. McGuire, Hop integrity in computer networks, *IEEE/ACM Trans. Netw.* 10 (3) (2002) 308–319.
- [39] A. Heffernan, Protection of BGP Sessions via the TCP MD5 Signature Option, (<https://tools.ietf.org/html/rfc2385>).
- [40] S. Kent, K. Seo, Security Architecture for the Internet Protocol, (<https://tools.ietf.org/html/rfc4301>).
- [41] J. Touch, Defending TCP Against Spoofing Attacks, (<https://tools.ietf.org/html/rfc4953>).
- [42] J. Touch, A. Mankin, R.P. Bonica, The TCP Authentication Option, (<https://tools.ietf.org/html/rfc5925>).
- [43] Y. Zhang, Z. Mao, J. Wang, Low-rate TCP-Targeted DoS attack disrupts internet routing, Symposium on Network and Distributed System Security (NDSS), (2007). San Diego, CA, USA
- [44] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S.F. Wu, L. Zhang, An analysis of BGP multiple origin AS (MOAS) conflicts, 1st SIGCOMM Workshop on Internet measurement, ACM, San Francisco, California, USA, 2001, pp. 31–35.
- [45] P. Sermpezis, V. Kotronis, A. Dainotti, X. Dimitropoulos, A Survey among Network Operators on BGP Prefix Hijacking, Technical Report, (2018).
- [46] A. Lutu, M. Bagnulo, C. Pelsser, K. Cho, R. Stanojevic, An analysis of the economic impact of strategic deaggregation, *Comput. Netw.* 81 (2015) 147–163.
- [47] X. Hu, Z.M. Mao, Accurate real-time identification of IP prefix hijacking, Symposium on Security and Privacy (S&P), IEEE, Oakland, CA, USA, 2007, pp. 3–17.
- [48] A. Cohen, Y. Gilad, A. Herzberg, M. Schapira, Jumpstarting BGP security with path-end validation, SIGCOMM Conference, ACM, Florianopolis, Brazil, 2016, pp. 342–355.
- [49] S. Goldberg, S. Halevi, A.D. Jaggard, V. Ramachandran, R.N. Wright, Rationality and traffic attraction: incentives for honest path announcements in BGP, SIGCOMM Conference on Data Communication, ACM, Seattle, WA, USA, 2008, pp. 267–278.
- [50] Q. Li, X. Zhang, X. Zhang, P. Su, Invalidating idealized BGP security proposals and countermeasures, *IEEE Trans. Dependable Secure Comput.* 12 (3) (2015) 298–311.
- [51] Y. Song, A. Venkataramani, L. Gao, Identifying and addressing reachability and policy attacks in secure BGP, *IEEE/ACM Trans. Netw.* 24 (5) (2016) 2969–2982.
- [52] M. Schuchard, E.Y. Vaserman, A. Mohaisen, D.F. Kune, N. Hopper, Y. Kim, Losing control of the internet: using the data plane to attack the control plane, Symposium on Network and Distributed System Security (NDSS), (2011). San Diego, CA, USA
- [53] J. Qiu, L. Gao, S. Ranjan, A. Nucci, Detecting bogus BGP route information: going beyond prefix hijacking, 3rd International Conference on Security and Privacy in Communications Networks and the Workshops (SecureComm), IEEE, Nice, France, 2007, pp. 381–390.
- [54] A. Toonk, BGP Leak Causing Internet Outages in Japan and Beyond, 2017, (<https://bgpmon.net/bgp-leak-causing-internet-outages-in-japan-and-beyond/>).
- [55] S. Murphy, BGP Security Vulnerabilities Analysis, (<https://www.ietf.org/rfc/rfc4272.txt>).

- [56] R. Perlman, Network Layer Protocols with Byzantine Robustness, Massachusetts Institute of Technology, Cambridge, 1988 Ph.D. thesis.
- [57] S. Bellovin, R. Bush, T.G. Griffin, J. Rexford, Slowing Routing Table Growth by Filtering Based on Address Allocation Policies, 2001, (<http://www.research.att.com/jrex>).
- [58] V. Khare, Q. Ju, B. Zhang, Concurrent prefix hijacks: occurrence and impacts, Internet Measurement Conference (IMC), ACM, Boston, Massachusetts, USA, 2012, pp. 29–36.
- [59] CIDR, (<http://www.cidr-report.org/>).
- [60] Practical Everyday BGP Filtering with AS_PATH Filters: Peer Locking, 2016, (https://www.nanog.org/sites/default/files/Snijders_Everyday_Practical_Bgp.pdf).
- [61] R. NCC, Routing Registry Consistency Check (RRCC), (<https://www.ripe.net/analyse/archived-projects/rrcc/>).
- [62] G. Siganos, M. Faloutsos, Analyzing BGP policies: methodology and tool, 23th International Conference on Computer Communications (INFOCOM), IEEE, Hong Kong, China, 2004, pp. 1640–1651.
- [63] R. NCC, RIPE Database Inconsistency Statistics, (<https://www.ripe.net/analyse/archived-projects/dbconstat/stats/ripe-database-inconsistency-statistics>).
- [64] E. yong Kim, L. Xiao, K. Nahrstedt, K. Park, Secure interdomain routing registry, IEEE Trans. Inf. Forensics Secur. 3 (2) (2008) 304–316.
- [65] T. Bates, R. Bush, T. Li, Y. Rekhter, DNS-based NLRI Origin AS Verification in BGP, (<https://tools.ietf.org/html/draft-bates-bgp4-nlri-orig-verif-00>).
- [66] J. Gersch, D. Massey, Reverse DNS Naming Convention for CIDR Address Blocks, (<https://tools.ietf.org/html/draft-gersch-dnsop-revdns-cidr-00>).
- [67] J. Gersch, D. Massey, ROVER: route origin verification using DNS, 22nd International Conference on Computer Communications and Networks (ICCCN), IEEE, Nassau, Bahamas, 2013.
- [68] G. Goodell, B. Aiello, T. Griffin, J. Ioannidis, A.R. Patrick McDaniel, Working around BGP: an incremental approach to improving security and accuracy of interdomain routing, Symposium on Network and Distributed System Security (NDSS), IEEE, San Diego, CA, USA, 2003.
- [69] A. Chen, A. Haeberlen, PRISM: private retrieval of the internet’s sensitive metadata, 8th USENIX Conference on Cyber Security Experimentation and Test, USENIX Association, Washington, DC, USA, 2015.
- [70] N. Feamster, H. Balakrishnan, J. Rexford, A. Shaikh, J. van der Merwe, The case for separating routing from routers, SIGCOMM Workshop on Future Directions in Network Architecture, ACM, Portland, Oregon, USA, 2004, pp. 5–12.
- [71] V. Kotronis, X. Dimitropoulos, B. Ager, Outsourcing the routing control logic: better internet routing based on SDN principles, 11th ACM Workshop on Hot Topics in Networks (HotNets), ACM, Redmond, Washington, USA, 2012, pp. 55–60.
- [72] V. Kotronis, A. Gämperli, X. Dimitropoulos, Routing centralization across domains via SDN: a model and emulation framework for BGP evolution, Comput. Netw. 92 (2) (2015) 227–239.
- [73] A. Gupta, L. Vanbever, M. Shahbaz, S.P. Donovan, B. Schlinker, N. Feamster, J. Rexford, S. Shenker, R. Clark, E. Katz-Bassett, SDX: a software defined internet exchange, SIGCOMM, ACM, Chicago, Illinois, USA, 2014, pp. 551–562.
- [74] A. Gupta, R. MacDavid, R. Birkner, M. Canini, N. Feamster, J. Rexford, L. Vanbever, An industrial-scale software defined internet exchange point, 13th USENIX Symposium on Networked Systems Design and Implementation (NSDI), USENIX Association, Santa Clara, CA, USA, 2016, pp. 1–14.
- [75] S. Kent, Securing the border gateway protocol: a status update, 7th IFIP-TC6 TC11 International Conference (CMS), Springer, Italy, 2003, pp. 40–53.
- [76] S. Kent, C. Lynn, J. Mikkelsen, K. Seo, Secure border gateway protocol (S-BGP) – real world performance and deployment issues, Symposium on Network and Distributed System Security (NDSS), IEEE, San Diego, CA, USA, 2000.
- [77] D.M. Nicol, S.W. Smith, M. Zhao, Evaluation of efficient security for BGP route announcements using parallel simulation, Simul. Model. Practice Theory 12 (3–4) (2004) 187–216.
- [78] J. Ng, Extensions to BGP to Support Secure Origin BGP (soBGP), (<https://tools.ietf.org/html/draft-ng-sobgp-bgp-extensions-02>).
- [79] P.C. van Oorschot, T. Wan, E. Kranakis, On interdomain routing security and pretty secure BGP (psBGP), ACM Trans. Inf. Syst. Secur. 10 (3) (2007).
- [80] T. Wan, E. Kranakis, P.C. van Oorschot, Pretty secure BGP (psBGP), Symposium on Network and Distributed System Security (NDSS), IEEE, San Diego, CA, USA, 2005.
- [81] Y.-C. Hu, D. McGrew, A. Perrig, B. Weis, D. Wendlandt, (R)Evolutionary bootstrapping of a global PKI for securing BGP, Fifth Workshop on Hot Topics in Networks (HotNets-V), (2006).
- [82] M. Lepinski, S. Kent, An Infrastructure to Support Secure Internet Routing, (<https://tools.ietf.org/html/rfc6480>).
- [83] G. Huston, G. Michaelson, Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs), (<https://tools.ietf.org/html/rfc6483>).
- [84] M. Lepinski, K. Sriram, BGPsec Protocol Specification, (<https://tools.ietf.org/html/draft-ietf-sidr-bgpsec-protocol-22>).
- [85] B. Weis, R. Gagliano, K. Patel, BGPsec Router Certificate Rollover, (<https://tools.ietf.org/html/draft-ietf-sidrops-bgpsec-rollover-01>).
- [86] Y.-J. Chi, R. Oliveira, L. Zhang, Cyclops: the AS-level connectivity observatory, ACM SIGCOMM Comput. Commun. Rev. 38 (5) (2008) 5–16.
- [87] K. Sriram, D. Montgomery, B. Dickson, K. Patel, A. Robachevsky, Methods for Detection and Mitigation of BGP Route Leaks, (<https://tools.ietf.org/html/draft-ietf-idr-route-leak-detection-mitigation-05>).
- [88] A. Cohen, Y. Gilad, A. Herzberg, M. Schapira, One hop for RPKI, one giant leap for BGP security, 14th Workshop on Hot Topics in Networks (HotNets), ACM, Philadelphia, PA, USA, 2015.
- [89] S. Goldberg, M. Schapira, P. Hummon, J. Rexford, How secure are secure interdomain routing protocols? SIGCOMM Conference, ACM, New Delhi, India, 2010, pp. 87–98.
- [90] G. Huston, Measures of Self-similarity of BGP updates and implications for securing BGP, 8th International Conference on Passive and Active Network Measurement, Springer, Louvain-la-Neuve, Belgium, 2007, pp. 1–10.
- [91] V.K. Sriram, D. Montgomery, Design and analysis of optimization algorithms to minimize cryptographic processing in BGP security protocols, Comput. Commun. 106 (2017) 75–85.
- [92] W. Aiello, J. Ioannidis, P. McDaniel, Origin authentication in interdomain routing, 10th Conference on Computer and Communications Security (CCS), ACM, Washington D.C., USA, 2003, pp. 165–178.
- [93] P. McDaniel, W. Aiello, K. Butler, J. Ioannidis, Origin authentication in interdomain routing, Comput. Netw. 50 (16) (2006) 2953–2980.
- [94] M. Zhao, S.W. Smith, D.M. Nicol, Aggregated path authentication for efficient BGP security, 12th Conference on Computer and Communications Security (CCS), ACM, Alexandria, VA, USA, 2005, pp. 128–138.
- [95] K. Brogle, S. Goldberg, L. Reyzin, Sequential aggregate signatures with lazy verification from trapdoor permutations, Inf. Comput. 239 (2014) 356–376.
- [96] A. Boldyreva, C. Gentry, A. O’Neill, D.H. Yum, Ordered Multisignatures and Identity-based Sequential Aggregate Signatures, with Applications to Secure Routing, 14th Conference on Computer and Communications Security (CCS), ACM, Alexandria, Virginia, USA, 2007, pp. 276–285.
- [97] A. Boldyreva, C. Gentry, A. O’Neill, D.H. Yum, New multiparty signature schemes for network routing applications, ACM Trans. Inf. Syst. Secur. 12 (1) (2008).
- [98] L.V. Mancini, A. Spognardi, C. Soriente, A. Villani, D. Vitali, Relieve internet routing security of public key infrastructure, 21st International Conference on Computer Communications and Networks (ICCCN), IEEE, Munich, Germany, 2012, pp. July–August.
- [99] Q. Li, M. Xu, J. Wu, X. Zhang, P.P.C. Lee, K. Xu, Enhancing the trust of internet routing with lightweight route attestation, IEEE Trans. Inf. Forensics Secur. 7 (2) (2012) 691–703.
- [100] Y. Xiang, X. Shi, J. Wu, Z. Wang, X. Yin, Sign what you really care about – secure BGP AS-paths efficiently, Comput. Netw. 57 (10) (2013) 2250–2265.
- [101] K. Butler, P. McDaniel, W. Aiello, Optimizing BGP security by exploiting path stability, 13th Conference on Computer and Communications Security (CCS), ACM, Alexandria, Virginia, USA, 2006, pp. 298–310.
- [102] Y.-C. Hu, A. Perrig, D.B. Johnson, Efficient security mechanisms for routing protocols, Symposium on Network and Distributed System Security (NDSS), IEEE, San Diego, CA, USA, 2003.
- [103] B. Raghavan, S. Panjwani, A. Mityagin, Analysis of the SPV secure routing protocol: weaknesses and lessons, ACM SIGCOMM Comput. Commun. Rev. 37 (2) (2007) 29–38.
- [104] H. Yin, B. Sheng, H. Wang, J. Pan, Keychain-based signatures for securing BGP, IEEE J. Sel. Areas Commun. 28 (8) (2010) 1308–1318.
- [105] B. Bruhadeshwar, S.S. Kulkarni, A.X. Liu, Symmetric key approaches to securing BGP – a little bit trust is enough, IEEE Trans. Parallel Distrib. Syst. 22 (9) (2011) 1536–1549.
- [106] D. Gupta, A. Segal, A. Panda, G. Segev, M. Schapira, J. Feigenbaum, J. Rexford, S. Shenker, A new approach to interdomain routing based on secure multi-party computation, 11th Workshop on Hot Topics in Networks (HotNets), ACM, Redmond, Washington, USA, 2012, pp. 37–42.
- [107] G. Asharov, D. Demmler, M. Schapira, T. Schneider, G. Segev, S. Shenker, M. Zohner, Privacy-preserving interdomain routing at internet scale, Privacy Enhancing Technol. (3) (2017) 1–21.
- [108] W. Henecka, M. Roughan, STRIP: privacy-preserving vector-based routing, 21st International Conference on Network Protocols (ICNP), IEEE, 2013.
- [109] A. Hari, T.V. Lakshman, The internet blockchain: a distributed, tamper-resistant transaction framework for the internet, 15th Workshop on Hot Topics in Networks (HotNets), ACM, Atlanta, GA, USA, 2016, pp. 204–210.
- [110] S. Goldberg, S. Halevi, A.D. Jaggard, V. Ramachandran, R.N. Wright, Rationality and traffic attraction: incentives for honest path announcements in BGP, ACM SIGCOMM Comput. Commun. Rev. 38 (4) (2008) 267–278.
- [111] Z.M. Mao, J. Rexford, J. Wang, R.H. Katz, Towards an accurate AS-level traceroute tool, SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, ACM, Karlsruhe, Germany, 2003, pp. 365–378.
- [112] V.N. Padmanabhan, D.R. Simon, Secure traceroute to detect faulty or malicious routing, ACM SIGCOMM Comput. Commun. Rev. 33 (1) (2003) 77–82.
- [113] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedmann, M. Latapy, C. Magnien, R. Teixeira, Avoiding traceroute anomalies with paris traceroute, Conference on Internet Measurement (IMC), ACM, Rio de Janeiro, Brazil, 2006, pp. 153–158.
- [114] I. Avramopoulos, J. Rexford, Stealth probing: efficient data-plane security for IP routing, Annual Conference on USENIX Annual Technical Conference, USENIX Association, Boston, MA, USA, 2006, pp. 25–31.
- [115] I. Avramopoulos, H. Kobayashi, R. Wang, A. Krishnamurthy, Highly secure and efficient routing, 23rd Annual Joint Conference of the Computer and Communications Societies (INFOCOM), IEEE, Hong Kong, China, 2004.
- [116] E.L. Wong, P. Balasubramanian, L. Alvisi, M.G. Gouda, V. Shmatikov, Truth in advertising: lightweight verification of route integrity, 26th Annual Symposium on Principles of Distributed Computing, ACM, Portland, Oregon, USA, 2007, pp. 147–156.
- [117] B. Bruhadeshwar, K. Kothapalli, M. Poornima, M. Divya, Routing protocol security using symmetric key based techniques, International Conference on Availability,

- Reliability and Security (ARES), IEEE, Fukuoka, Japan, 2009.
- [118] T.H.-J. Kim, C. Basescu, L. Jia, S.B. Lee, Y.-C. Hu, A. Perrig, Lightweight source authentication and path validation, SIGCOMM Conference, ACM, Chicago, Illinois, USA, 2014, pp. 271–282.
- [119] X. Liu, X. Yang, D. Wetherall, T. Anderson, Efficient and secure source authentication with packet passports, 2nd Conference on Steps to Reducing Unwanted Traffic on the Internet (SRUTI), USENIX Association, San Jose, CA, USA, 2006.
- [120] X. Liu, A. Li, X. Yang, D. Wetherall, Passport: secure and adoptable source authentication, 5th Symposium on Networked Systems Design and Implementation (NDSI), USENIX Association, San Francisco, CA, USA, 2008, pp. 365–378.
- [121] J. Naoaus, M. Walfish, A. Nicolosi, D. Mazires, M. Miller, A. Seehra, Verifying and enforcing network paths with icing, 7th Conference on Emerging Networking Experiments and Technologies (CoNEXT), ACM, Tokyo, Japan, 2011.
- [122] C. Pappas, R.M. Reischuk, A. Perrig, FAIR: forwarding accountability for internet reputability, Technical Report, (2016).
- [123] L. Subramanian, V. Roth, I. Stoica, S. Shenker, R.H. Katz, Listen and whisper: security mechanisms for BGP, 1st Conference on Symposium on Networked Systems Design and Implementation (NDSI), USENIX, San Francisco, CA, USA, 2004.
- [124] A.T. Mizrak, Y.-C. Cheng, K. Marzullo, S. Savage, Fatih: detecting and isolating malicious routers, International Conference on Dependable Systems and Networks (DSN), IEEE, Yokohama, Japan, 2005, pp. 538–547.
- [125] A.T. Mizrak, Y.-C. Cheng, K. Marzullo, S. Savage, Detecting and isolating malicious routers, IEEE Trans. Dependable Secure Comput. 3 (3) (2006) 230–244.
- [126] K. Argyraki, P. Maniatis, O. Irazak, S. Ashish, S. Shenker, Loss and delay accountability for the internet, International Conference on Network Protocols (ICNP), IEEE, Beijing, China, 2007, pp. 194–205.
- [127] C. Basescu, Y.-H. Lin, H. Zhang, A. Perrig, High-speed inter-domain fault localization, Symposium on Security and Privacy (S&P), San Jose, CA, USA, IEEE, 2016, pp. 859–877.
- [128] X. Zhang, C. Lan, A. Perrig, Secure and scalable fault localization under dynamic traffic patterns, Symposium on Security and Privacy (S&P), IEEE, San Francisco, CA, USA, 2012, pp. 317–331.
- [129] X. Zhang, Z. Zhou, H.-C. Hsiao, T.H.-J. Kim, A. Perrig, P. Tague, Shortmac: efficient data-plane fault localization, Symposium on Network and Distributed System Security (NDSS), (2012). San Diego, CA, USA
- [130] D. Zhu, M. Gritter, D.R. Cheriton, Feedback based routing, ACM SIGCOMM Comput. Commun. Rev. 33 (1) (2003) 71–76.
- [131] D. Wendlandt, I. Avramopoulos, D.G. Andersen, J. Rexford, Don't secure routing protocols, secure data delivery, 5th Workshop on Hot Topics in Networks (HotNets), ACM, Irvine, CA, 2006.
- [132] H. Chan, D. Dash, H. Zhang, Modeling adoptability of secure BGP protocols, Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM), ACM, Pisa, Italy, 2006, pp. 279–290.
- [133] I. Avramopoulos, M. Suchara, J. Rexford, How Small Groups Can Secure Interdomain Routing, (Technical Report), (2007).
- [134] H. Ballani, P. Francis, X. Zhang, A study of prefix hijacking and interception in the internet, Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, ACM, Kyoto, Japan, 2007, pp. 265–276.
- [135] P. Gill, M. Schapira, S. Goldberg, A survey of interdomain routing policies, ACM SIGCOMM Comput. Commun. Rev. 44 (2014) 28–34.
- [136] R. Bush, R. Austein, **The Resource Public Key Infrastructure (RPKI) to Router Protocol**, (<https://tools.ietf.org/html/rfc6810>).
- [137] M. Wählisch, R. Schmidt, T.C. Schmidt, O. Maennel, S. Uhlig, G. Tyson, RiPKI: the tragic story of RPKI deployment in the web ecosystem, 14th Workshop on Hot Topics in Networks (HotNets), ACM, Philadelphia, PA, USA, 2015.
- [138] Y. Gilad, A. Cohen, A. Herzberg, M. Schapira, H. Shulman, Are we there yet? on RPKI's deployment and security, Symposium on Network and Distributed System Security (NDSS), Internet Society, San Diego, CA, USA, 2017.
- [139] D. Cooper, E. Heilman, K. Brogle, L. Reyzin, S. Goldberg, On the risk of misbehaving RPKI authorities, 12th Workshop on Hot Topics in Networks (HotNets), ACM, College Park, Maryland, 2013.
- [140] E. Heilman, D. Cooper, L. Reyzin, S. Goldberg, From the consent of the routed: improving the transparency of the RPKI, SIGCOMM Conference, ACM, Chicago, Illinois, USA, 2014, pp. 51–62.
- [141] E. Osterweil, T. Manderson, R. White, Sizing Estimates for a Fully Deployed RPKI, 2012, (<https://techreports.verisignlabs.com/docs/tr-1120005-2.pdf>).
- [142] S. Kent, K. Sriram, RPKI rsync Download Delay Modeling, 2013, (<https://www.ietf.org/proceedings/86/slides/slides-86-sidr-1.pdf>).
- [143] T. Bruijnzeels, O. Muravskiy, B. Weber, RPKI Repository Analysis and Requirements, 2013, (<https://tools.ietf.org/pdf/draft-tbruijnzeels-sidr-repo-analysis-00.pdf>).
- [144] S. Goldberg, M. Schapira, P. Hummon, J. Rexford, How secure are secure interdomain routing protocols? Comput. Netw. 70 (2014) 260–287.
- [145] A. Boldyreva, R. Lychev, Provable security of S-BGP and other path vector protocols: model, analysis and extensions, Conference on Computer and Communications Security (CCS), ACM, Raleigh, North Carolina, USA, 2012, pp. 541–552.
- [146] S.Y. Qiu, F. Monrose, A. Terzis, P.D. McDaniel, Efficient techniques for detecting false origin advertisements in inter-domain routing, 2nd Workshop on Secure Network Protocols, IEEE, Santa Barbara, CA, USA, 2006, pp. 12–19.
- [147] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, L. Zhang, PHAS: a prefix hijack alert system, 15th Conference on USENIX Security Symposium, USENIX Association, Vancouver, B.C., Canada, 2006, pp. 153–166.
- [148] G. Siganos, M. Faloutsos, Neighborhood watch for internet routing: can we improve the robustness of internet routing today? 26th International Conference on Computer Communications (INFOCOM), IEEE, Barcelona, Spain, 2007, pp. 1271–1279.
- [149] R. Hiran, N. Carlsson, N. Shahmehri, PrefiSec: a distributed alliance framework for collaborative BGP monitoring and prefix-based security, Workshop on Information Sharing & Collaborative Security, ACM, Scottsdale, Arizona, USA, 2014, pp. 3–12.
- [150] R. Hiran, N. Carlsson, N. Shahmehri, Collaborative framework for protection against attacks targeting BGP and edge networks, Comput. Netw. 122 (2017) 120–137.
- [151] A. Haeberlen, I. Avramopoulos, J. Rexford, P. Druschel, NetReview: detecting when interdomain routing goes wrong, 6th Symposium on Networked Systems Design and Implementation (NDSI), USENIX Association, Boston, Massachusetts, 2009, pp. 437–452.
- [152] A.J.T. Gurney, A. Haeberlen, W. Zhou, M. Sherr, B.T. Loo, Having your cake and eating it too: routing security with privacy protections, 10th Workshop on Hot Topics in Networks (HotNets), ACM, Cambridge, Massachusetts, USA, 2011.
- [153] S. Nagaraja, V. Jalaparti, M. Caesar, N. Borisov, P3CA: private anomaly detection across ISP networks, 11th International Symposium on Privacy Enhancing Technologies, Springer, Waterloo, ON, Canada, 2011, pp. 38–56.
- [154] M. Zhao, W. Zhou, A.J.T. Gurney, A. Haeberlen, M. Sherr, B.T. Loo, Private and verifiable interdomain routing decisions, Applications, Technologies, Architectures, and Protocols for Computer Communication, ACM, Helsinki, Finland, 2012, pp. 383–394.
- [155] M. Zhao, W. Zhou, A.J.T. Gurney, A. Haeberlen, M. Sherr, B.T. Loo, Private and verifiable interdomain routing decisions, IEEE/ACM Trans. Netw. 24 (2) (2016) 1011–1024.
- [156] N.M. Al-Rousan, L. Trajković, Machine learning models for classification of BGP anomalies, 13th International Conference on High Performance Switching and Routing (HPSR), IEEE, Belgrade, Serbia, 2012, pp. 103–108.
- [157] Y. Guo, H. Duan, J. Chen, F. Miao, MAF-SAM: An effective method to perceive data plane threats of inter domain routing system, Comput. Netw. 110 (2016) 69–78.
- [158] M. Zhang, J. Li, S. Brooks, 1-seismograph: observing, measuring, and analyzing internet earthquakes, IEEE/ACM Trans. Netw. 25 (6) (2017) 3411–3426.
- [159] A. Lutu, M. Bagnulo, C. Pelsser, O. Maennel, J. Cid-Sueiro, The BGP visibility toolkit: detecting anomalous internet routing behavior, IEEE/ACM Trans. Netw. 24 (2) (2016) 1237–1250.
- [160] S. Li, H. Duan, Z. Wang, X. Li, Route leaks identification by detecting routing loops, Security and Privacy in Communication Networks, (2015), pp. 313–329.
- [161] J. Chang, K.K. Venkatasubramanian, A.G. West, S. Kannan, I. Lee, B.T. Loo, O. Sokolsky, AS-CRED: Reputation and alert service for interdomain routing, IEEE Syst. J. 7 (3) (2013) 396–409.
- [162] S. Su, B. Zhang, B. Fang, Online detection of concurrent prefix hijacks, International Conference on Security and Privacy in Communication Networks (SecureComm), Springer, Beijing, China, 2014, pp. 69–83.
- [163] S. Papadopoulos, K. Moustakas, D. Tzovaras, BGPViewer: using graph representations to explore BGP routing changes, 18th International Conference on Digital Signal Processing (DSP), IEEE, Fira, Greece, 2013.
- [164] S. Papadopoulos, G. Theodoridis, D. Tzovaras, BGPfuser: using visual feature fusion for the detection and attribution of BGP anomalies, 10th Workshop on Visualization for Cyber Security (VizSec), ACM, Atlanta, Georgia, USA, 2013, pp. 57–64.
- [165] M. Syamkumar, R. Durairajan, P. Barford, Bigfoot: a geo-based visualization methodology for detecting BGP threats, Symposium on Visualization for Cyber Security (VizSec), IEEE, Baltimore, MD, USA, 2016.
- [166] C. Zheng, L. Ji, D. Pei, J. Wang, P. Francis, A light-weight distributed scheme for detecting IP prefix hijacks in real-time, Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM), ACM, Kyoto, Japan, 2007, pp. 277–288.
- [167] L. Quan, J. Heidemann, Y. Pradkin, Trinocular: understanding internet reliability through adaptive probing, SIGCOMM, ACM, Hong Kong, China, 2013, pp. 255–266.
- [168] R. Hiran, N. Carlsson, N. Shahmehri, Crowd-based detection of routing anomalies on the internet, Conference on Communications and Network Security (CNS), IEEE, Florence, Italy, 2015, pp. 388–396.
- [169] K. Balu, M.L. Pardal, M. Correia, DARSHANA: detecting route hijacking for communication confidentiality, 15th International Symposium on Network Computing and Applications (NCA), IEEE, Cambridge, MA, USA, 2016, pp. 52–59.
- [170] Z. Zhang, Y. Zhang, Y.C. Hu, Z.M. Mao, R. Bush, iSPY: detecting IP prefix hijacking on my own, SIGCOMM Conference on Data Communication, ACM, Seattle, WA, USA, 2008, pp. 327–338.
- [171] M. Djatmiko, D. Schatzmann, X. Dimitropoulos, A. Friedman, R. Boreli, Federated flow-based approach for privacy preserving connectivity tracking, 9th ACM conference on Emerging networking experiments and technologies (CoNEXT), ACM, Santa Barbara, CA, USA, 2013, pp. 429–440.
- [172] S. Goldberg, D. Xiao, E. Tromer, B. Barak, J. Rexford, Path-quality monitoring in the presence of adversaries, International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS), ACM, Annapolis, MD, USA, 2008, pp. 193–204.
- [173] S. Goldberg, D. Xiao, E. Tromer, B. Barak, J. Rexford, Path-quality monitoring in the presence of adversaries: the secure sketch protocols, IEEE/ACM Trans. Netw. 23 (6) (2015) 1729–1741.
- [174] X. Zhang, A. Jain, A. Perrig, Packet-dropping adversary identification for data plane security, CoNEXT Conference, ACM, Madrid, Spain, 2008.
- [175] F. Fischer, J. Fuchs, P.-A. Vervier, F. Mansmann, O. Thonnard, VisTracer: a visual analytics tool to investigate routing anomalies in traceroutes, 9th International Symposium on Visualization for Cyber Security (VizSec), ACM, Seattle, Washington, USA, 2012, pp. 80–87.

- [176] Y. Xiang, Z. Wang, X. Yin, J. Wu, Argus: an accurate and agile system to detecting IP prefix hijacking, 19th International Conference on Network Protocols (ICNP), IEEE, Vancouver, BC, Canada, 2011.
- [177] X. Shi, Y. Xiang, Z. Wang, X. Yin, J. Wu, Detecting prefix hijackings in the internet with argus, Internet Measurement Conference (IMC), ACM, Boston, Massachusetts, USA, 2012, pp. 15–28.
- [178] J. Schlamp, R. Holz, O. Gasser, A. Korsten, Q. Jacquemart, G. Carle, E.W. Biersack, Investigating the Nature of Routing Anomalies: Closing in on Subprefix Hijacking Attacks, 7th International Workshop on Traffic Monitoring and Analysis, Springer, Barcelona, Spain, 2015, pp. 173–187.
- [179] J. Schlamp, R. Holz, Q. Jacquemart, G. Carle, E.W. Biersack, HEAP: Reliable assessment of BGP hijacking attacks, IEEE J. Selected Areas Commun. 34 (6) (2016) 1849–1861.
- [180] M.S. Siddiqui, D. Montero, R. Serral-Gracià, X. Masip-Bruin, W. Ramirez, Route leak detection using real-time analytics on local BGP information, Global Communications Conference (GLOBECOM), IEEE, Austin, TX, USA, 2014.
- [181] M.S. Siddiqui, D. Montero, R. Serral-Gracià, M. Yannuzzi, Self-reliant detection of route leaks in inter-domain routing, Comput. Netw. 82 (2015) 135–155.
- [182] A. Feldmann, O. Maennel, Z.M. Mao, A. Berger, B. Maggs, Locating internet routing instabilities, SIGCOMM, ACM, Portland, Oregon, USA, 2004, pp. 205–218.
- [183] U. Javed, I. Cunha, D. Choffnes, E. Katz-Bassett, T. Anderson, A. Krishnamurthy, PoiRoot: investigating the root cause of interdomain path changes, Conference on SIGCOMM, ACM, Hong Kong, China, 2013, pp. 183–194.
- [184] E. Katz-Bassett, H.V. Madhyastha, J.P. John, A. Krishnamurthy, T. Anderson, Studying black holes in the internet with hubble, 5th USENIX Symposium on Networked Systems Design and Implementation (NSDI), USENIX, San Francisco, CA, USA, 2008, pp. 247–262.
- [185] E. Katz-Bassett, C. Scott, D.R. Choffnes, Ítalo Cunha, V. Valancius, N. Feamster, H.V. Madhyastha, T. Anderson, A. Krishnamurthy, LIFEGUARD: practical repair of persistent route failures, SIGCOMM, ACM, Helsinki, Finland, 2012, pp. 395–406.
- [186] T. Qiu, L. Ji, D. Pei, J. Wang, J. Xu, H. Ballani, Locating prefix hijackers using LOCK, 18th Conference on USENIX Security Symposium, USENIX Association, Montreal, Canada, 2009, pp. 135–150.
- [187] L. Wang, X. Zhao, D. Pei, R. Bush, D. Massey, L. Zhang, Protecting BGP routes to top-level DNS servers, IEEE Trans. Parallel Distrib. Syst. 14 (9) (2003) 851–860.
- [188] Z. Zhang, Y. Zhang, Y.C. Hu, Z.M. Mao, Practical defenses against BGP prefix hijacking, CoNEXT Conference, ACM, New York, NY, USA, 2007.
- [189] J. Karlin, S. Forrest, J. Rexford, Pretty good BGP: improving BGP by cautiously adopting routes, International Conference on Network Protocols (ICNP), IEEE, Santa Barbara, CA, USA, 2006, pp. 290–299.
- [190] J. Karlin, S. Forrest, J. Rexford, Autonomous security for autonomous systems, J. Comput. Netw. 15 (52) (2008) 2908–2923.
- [191] M. Zhang, B. Liu, B. Zhang, Safeguarding data delivery by decoupling path propagation and adoption, INFOCOM, IEEE, San Diego, CA, USA, 2010.
- [192] T. Holterbach, S. Vissicchio, A. Dainotti, L. Vanbever, SWIFT: predictive fast route, SIGCOMM, ACM, Los Angeles, CA, USA, 2017, pp. 460–473.
- [193] P. Sermpezis, V. Kotronis, P. Gigis, X. Dimitropoulos, D. Cicalese, A. King, A. Dainotti, ARTEMIS: neutralizing BGP hijacking within a minute, Technical Report, (2018).
- [194] T. Qiu, L. Ji, D. Pei, J. Wang, J. Xu, TowerDefense: deployment strategies for battling against IP prefix hijacking, 18th International Conference on Network Protocols (ICNP), IEEE, Kyoto, Japan, 2010.
- [195] P. Gill, M. Schapira, S. Goldberg, Modeling on quicksand: dealing with the scarcity of ground truth in interdomain routing data, ACM SIGCOMM Comput. Commun. Rev. 42 (1) (2012) 40–46.
- [196] R. Anwar, H. Niaz, D. Choffnes, Ítalo Cunha, P. Gill, E. Katz-Bassett, Investigating interdomain routing policies in the wild, Internet Measurement Conference (IMC), ACM, Tokyo, Japan, 2015, pp. 71–77.
- [197] M. Roughan, W. Willinger, O. Maennel, D. Perouli, R. Bush, 10 Lessons from 10 years of measuring and modeling the internet’s autonomous systems, IEEE J. Selected Areas Commun. 29 (9) (2011) 1810–1821.
- [198] C. Dietzel, A. Feldmann, T. King, Blackholing at IXPs: on the effectiveness of DDoS mitigation in the wild, 17th International Conference on Passive and Active Measurement (PAM), Springer, Heraklion, Greece, 2016, pp. 319–332.
- [199] V. Giotsas, P. Richter, G. Smaragdakis, A. Feldmann, C. Dietzel, A. Berger, Inferring BGP blackholing activity in the internet, Internet Measurement Conference (IMC), ACM, London, United Kingdom, 2017, pp. 1–14.